arXiv:1411.1459v2 [math.OC] 21 Mar 2015

Technical Report
Nov. 5, 2014 (revised Mar. 2015)

On Convergence of Value Iteration for a Class of Total Cost Markov Decision Processes*

Huizhen Yu†

Abstract

We consider a general class of total cost Markov decision processes (MDP) in which the one-stage costs can have arbitrary signs, but the sum of the negative parts of the one-stage costs is finite for all policies and all initial states. This class, which we refer to as the General Convergence (GC for short) total cost model, contains several special classes of problems, e.g., positive costs problems, bounded negative costs problems, and discounted problems with unbounded one-stage costs. We study the convergence of value iteration for the GC model, in the Borel MDP framework with universally measurable policies. Our main results include: (i) convergence of value iteration when starting from certain functions above the optimal cost function; (ii) convergence of transfinite value iteration starting from zero, in the special case where the optimal cost function is nonnegative; and (iii) partial convergence of value iteration starting from zero, for a subset of initial states.

These results extend several previously known results about the convergence of value iteration for either positive costs problems or GC total cost problems. In particular, the first result on convergence of value iteration from above extends a theorem of van der Wal for the GC model. The second result relates to Maitra and Sudderth's analysis of transfinite value iteration for the positive costs model. It suggests connections between the two total cost models when the optimal cost function is nonnegative, and it leads to additional results on the convergence of ordinary non-transfinite value iteration, with a suitably defined dynamic programming operator, for finite state or finite control GC problems. The third result on partial convergence of value iteration is motivated by Whittle's bridging condition for the positive costs model, and provides a novel extension of the bridging condition to the GC model, where there are no sign constraints on the costs.

Key words: discrete-time stochastic optimal control; Markov decision processes; infinite spaces; dynamic programming; value iteration; convergence

*This research is in part supported by a grant from Alberta Innovates – Technology Futures.
†Department of Computing Science, University of Alberta. Email: janey.hzyu@gmail.com

Contents

1 Introduction **3**

2 Background **4**
2.1 Borel-Spaces MDP and the Total Cost Model (GC) 4
2.2 Some Optimality Properties of the (GC) Model 7
2.3 The Question about Value Iteration 8

3 Convergence of Value Iteration from Above **9**

4 Convergence of Value Iteration starting from Zero **12**
4.1 Special Cases where J^* is Nonpositive or Nonnegative 13
4.2 The General Case 17
4.2.1 Preliminaries 17
4.2.2 Convergence Results 18
4.2.3 Implications and Discussion 20

5 Illustrative Examples **23**

6 Concluding Remarks **25**

References **26**

Appendices **28**

A Some Basic Properties of the (GC) Model in the Universal Measurability Framework **28**

B Optimality Properties of the (GC) Model in the Universal Measurability Framework **29**
B.1 A Deterministic Model and its Relation to the Original Problem 30
B.2 Proof Arguments for Theorems 2.1, 2.2 32
B.2.1 J^* is Lower Semi-analytic 32
B.2.2 Existence of an ϵ-Optimal Semi-Markov Policy 33
B.2.3 The Optimality Equation $J^* = T(J^*)$ 33
B.2.4 Existence of an ϵ-Optimal Nonrandomized Semi-Markov Policy 34
B.2.5 Existence of ϵ-Optimal Markov or Optimal Stationary Policies when $J^* \geq 0$. 35

C About Transfinite Value Iteration and an Important Lemma for Proving Theorem 4.1 **35**
C.1 Ordinals and Transfinite Value Iteration 36
C.2 A Lemma for Convergence of Transfinite Value Iteration 36

1 Introduction

In this paper we study convergence properties of value iteration for a class of Markov decision processes (MDP) under the undiscounted total cost criterion. Specifically, we consider problems in which the one-stage costs can take both positive and negative values, but we assume that under any policy and for any starting state, the expected total sum, over the infinite horizon, of the negative parts of the one-stage cost is finite. Following the extensive survey on this class of MDP (Feinberg [15]), we shall refer to this total cost model as General Convergence total cost model (GC for short). It contains several classes of special models, in particular, the positive costs model (P), where all the one-stage costs are nonnegative, the negative costs model (N) where all the one-stage costs are nonpositive and the optimal costs are finite, as well as the bounded negative costs model [28, Chap. 7.2] and a subset of the discounted models with unbounded costs (UD).

It is known that for the (GC) model, value iteration starting from the constant function zero need not converge to the optimal cost function, surprisingly, even for finite state and control problems (see van der Wal [37, Example 3.2], Feinberg [15, Example 6.10]). This is also true for the positive costs model (P) (although convergence is ensured for (P) if each state has only a finite number of feasible controls [2, Prop. 9.18] or more generally, if certain compactness conditions ([1], [2, Prop. 9.17]) or semicontinuity and compactness model assumptions [29] are satisfied). In the (P) case, Maitra and Sudderth established the convergence of transfinite value iteration [24], and Whittle formulated a sufficient condition, called the bridging condition, for the convergence of value iteration starting from zero as well as from any function between zero and a multiple of the optimal cost function [40] (see also [19, 28]). Recently, motivated by these earlier results, Yu and Bertsekas showed that for the (P) case, value iteration always converges from above, when it starts within the set of functions that are above the optimal cost function and yet bounded by a multiple of the optimal cost function [41]. In this paper, we extend these results for the (P) case to the total cost problems in the (GC) model, where there are no sign constraints on the one-stage costs.

The main contributions of this paper are as follows:

(i) We prove that value iteration always converges from above if it starts within a set of functions that are above the optimal cost function and yet bounded by a multiple of a certain nonnegative function. This nonnegative function is given by $J^{*+}+J^{*-}$, where J^{*+}, $J^{*-} \geq 0$ are the optimal cost functions with respect to the positive part and negative part of the one-stage cost function, respectively (see Theorem 3.1). This result extends a theorem of van der Wal [37, Theorem 3.7] about the convergence of value iteration for (GC).

(ii) We prove that if the optimal cost function is nonnegative, then when suitably defined for (GC), transfinite value iteration starting from zero converges to the optimal cost function (see Theorem 4.1). This result is analogous to the transfinite value iteration result of Maitra and Sudderth for positive costs problems [24]. We also show that if the optimal cost function is nonnegative, then (a) between zero and the optimal cost function, the dynamic programming operator has no fixed points other than the optimal cost function itself (see Cor. 4.1); and (b) convergence of ordinary non-transfinite value iteration, with a suitably defined operator, can be ensured for finite control problems (see Prop. 4.3), as well as for finite state problems that have finite optimal costs (see Prop. 4.2).

(iii) We formulate a novel sufficient condition, in the style of Whittle's bridging condition [40], for *partial convergence* of value iteration in (GC) (see Theorems 4.2, 4.3 and Remark 4.1). It characterizes the convergence of value iteration starting from zero, for a *subset* of states.

Besides the closely related early works mentioned above, we will discuss later, with more precise mathematical terms, some other related works in connection with our results in Sections 2.3 and 4.2.3.

To present our analyses of the (GC) model for general state and control spaces, we shall consider Borel-spaces MDP in the universal measurability framework (Shreve and Bertsekas [33], Bertsekas

and Shreve [2, Part II]). We shall derive various optimality and convergence properties for the (GC) model in this framework. Most of the ideas in our analyses of value iteration are, however, independent of the measurability concepts, so we will just summarize the universal measurability framework, selectively note the properties of analytic sets and universally measurable policies needed in our proofs, and refer to the original sources for some of the details. (As to the deep mathematical reason for adopting this framework, we refer the readers to the articles [7, 32] and the monograph [2].) To our knowledge, the (GC) model has not been thoroughly studied in the universal measurability framework, so some of the optimality properties of the (GC) model that we present in this paper can also be of independent interest.

While our focus will be on value iteration, it is worth noting that generally, for undiscounted total cost MDP, it is not easy to find near-optimal policies or optimal ones such as optimal stationary policies (if they exist), even if the optimal cost function is available. For positive or negative costs models as well as for the (GC) model, the existence proofs for near-optimal policies are based on constructing such a policy, using the optimal cost function in addition to some other information (cf. the proof of [2, Prop. 9.20] and the proof of Theorem 2.2(a) given in Appendix B of this paper). Thus in theory, methods of finding a near-optimal policy are available, although they may not be computationally tractable. As to finding an optimal stationary policy, the standard policy iteration method is known to have convergence issues even when such a policy exists (see e.g., [28, Chap. 7]); recently, in a broader context of constrained total cost MDP, a new method based on linear programming has been established by [11, 12] for semicontinuous models under certain conditions.

This paper is organized as follows. In Section 2 we define formally the class of total cost problems in (GC), using the universal measurability framework, and we summarize some of the fundamental optimality properties of this model. In Section 3 we establish the convergence of value iteration starting within a certain set of functions above the optimal cost function. In Section 4 we consider value iteration starting from the constant function zero. We first show, in Section 4.1, that if the optimal cost function is nonnegative, then the (GC) model exhibits some properties similar to those in the positive model (P). These include the fixed point properties of the dynamic programming operator, the convergence of transfinite value iteration, and the convergence of ordinary non-transfinite value iteration for finite state or finite control problems. We then present in Section 4.2 our extension of Whittle's bridging condition. We also discuss some direct consequences of this result, including one that relates to the sufficient conditions given earlier by van Hee et al. [38], and by Hernández-Lerma and Lasserre [20], for value iteration. In Section 5 we use examples to illustrate the convergence results obtained in this paper. In Section 6 we conclude the paper with several remarks on open questions for future research. Appendices A-C collect some proofs or outlines of proofs that require an in-depth use of the Borel MDP and the universal measurability framework.

2 Background

2.1 Borel-Spaces MDP and the Total Cost Model (GC)

We consider Borel-spaces MDP in the universal measurability framework [2, Part II]. Let S and C denote the *state space* and the *control space*, respectively. We assume that they are Borel spaces.[1] We write x for a state in S and u for a control in C. Associated with each state x is a nonempty set $U(x) \subset C$, consisting of *feasible controls* at x. The set-valued function $U : x \mapsto U(x)$ thus specifies the *control constraint*, and its graph $\Gamma = \{(x, u) \mid x \in S, u \in U(x)\}$ is assumed to be an analytic subset of $S \times C$.[2] A given function $g : \Gamma \to [-\infty, +\infty]$ specifies the *one-stage cost* $g(x, u)$ for each

[1] A Borel space is a Borel subset of some Polish space (a complete separable metric space).

[2] There are several equivalent definitions of analytic sets in a Polish space (see [2, Prop. 7.41], [10, Chap. 13.2]). The empty set is analytic. A nonempty analytic set is by definition the image of some Borel set in some Polish space, under a Borel measurable mapping. In particular, if X, Y are two Polish spaces and B is a Borel subset of $X \times Y$, the projection of B into X is an analytic set in X.

state and feasible control pair $(x, u) \in \Gamma$. We assume that the function g is lower semi-analytic.[3] At state x, if we apply a control $u \in U(x)$, a one-stage cost $g(x, u)$ is incurred, and the system then moves to the next state x' according to a given stochastic kernel $q(dx' \mid x, u)$. More precisely, the kernel q for *state transition* is defined as follows. For a Borel space Y, we denote by $\mathcal{P}(Y)$ the set of probability measures on the Borel σ-algebra on Y. The state transition kernel $q(dx' \mid x, u)$ maps each $(x, u) \in S \times C$ to a probability measure in $\mathcal{P}(S)$, and furthermore, q is assumed to be Borel measurable, meaning that the mapping $q : (x, u) \mapsto q(\cdot \mid x, u) \in \mathcal{P}(S)$ is Borel measurable (where the space $\mathcal{P}(S)$ is equipped with the weak topology).

For $k \geq 0$, we denote by (x_k, u_k) the state and control pair at time k. We define a *policy* for the above MDP to be a sequence of stochastic kernels, $\pi = (\mu_0, \mu_1, \ldots)$, where for each $k \geq 0$, μ_k is a universally measurable stochastic kernel[4] on C given $(S \times C)^k \times S$ such that

$$\mu_k(U(x_k) \mid x_0, x_1, \ldots, x_k) = 1, \qquad \forall\, (x_0, x_1, \ldots, x_k) \in (S \times C)^k \times S.$$

We denote by Π the set of all policies. For each initial state $x \in S$, a policy π induces a stochastic process $(x_0, u_0, x_1, u_1, \ldots)$ on the space $(S \times C)^\infty$ (endowed with the universal σ-algebra), with $x_0 = x$. The probability of this process is determined uniquely by the collection of the stochastic kernels comprising π, together with the state transition kernel $q(dx' \mid x, u)$ [2, Prop. 7.45]. We denote by $\mathbb{E}_x^\pi$ expectation with respect to this induced probability measure on $(S \times C)^\infty$.

We define structured families of policies in the standard way. A policy is *nonrandomized* if for each k and each $(x_0, u_0, \ldots, x_k)$, $\mu_k(du_k \mid x_0, x_1, \ldots, x_k)$ assigns probability one to a single control in $U(x_k)$. A policy is *semi-Markov* (resp. *Markov*) if for each k, the function $\mu_k : (x_0, u_0, \ldots, x_k) \mapsto \mu_k(\cdot \mid x_0, u_0, \ldots, x_k)$ depends only on (x_0, x_k) (resp. x_k). A Markov policy of the form $\pi = (\mu, \mu, \ldots)$ is called a *stationary* policy, in which case we simply write μ for that policy.

We mention that since the set Γ of state and feasible control pairs is analytic, a nonrandomized stationary policy exists by the Jankov-von Neumann selection theorem [2, Prop. 7.49], and therefore, the set Π of policies is nonempty.[5] We refer to the monograph [2, Part II] for a full introduction of the universal measurability framework, including the mathematical background for the preceding definitions, the properties of lower semi-analytic functions and universally measurable policies, and measurable selection theorems upon which the framework is built.[6]

The class (GC) of total cost problems is defined by a finiteness condition on the sums of the negative parts of the one-stage costs under any policy. More specifically, let us write g in terms of its positive part g_+ and negative part g_- as:

$$g = g_+ - g_-, \qquad \text{where} \quad g_+ = \max\{g, 0\}, \quad g_- = \max\{-g, 0\}.$$

For any policy $\pi \in \Pi$, define

$$J_\pi^+(x) = \mathbb{E}_x^\pi \left\{ \sum_{k=0}^\infty g_+(x_k, u_k) \right\}, \qquad J_\pi^-(x) = \mathbb{E}_x^\pi \left\{ \sum_{k=0}^\infty g_-(x_k, u_k) \right\}, \quad x \in S,$$

where the expectations are well-defined by the monotone convergence theorem.

[3] A function $f : D \to [-\infty, +\infty]$ is lower semi-analytic if and only if its domain D is an analytic set and its level sets $\{y \in D \mid f(y) < c\}$ are analytic for all real c [2, Def. 7.21]. Equivalently, f is lower semi-analytic if and only if its epigraph $\{(y, c) \mid f(y) \leq c,\, y \in D,\, c \in (-\infty, +\infty)\}$ is analytic (see [2, p. 186], [24]).

[4] For two Borel spaces Y, Z, a universally measurable stochastic kernel on Y given Z is by definition a universally measurable mapping from Z to the space $\mathcal{P}(Y)$. Here $\mathcal{P}(Y)$ is equipped with the weak topology, which makes $\mathcal{P}(Y)$ a Borel space [2, Cor. 7.25.1], and the σ-algebra on Z is the universal σ-algebra (see e.g., [2, Def. 7.18] and [10, Chap. 3.3] for the completion of measures and the universal σ-algebra).

[5] By comparison, if Γ is assumed to be Borel measurable instead and only Borel measurable policies are allowed, then, without additional assumptions, the set of policies can be empty [6].

[6] For a quick overview of some of these subjects, we refer the readers to the articles [7, 24, 32] and the recent paper [41, Section 2]. For a general introduction to Borel MDP with Borel measurable policies, see the book [13]; for analytic sets, see also [10, Chap. 13.2] and [25, Chap. 6.2].

Definition 2.1 (The (GC) total cost model)**.** An MDP given above is in the *(GC)* class if for all policies $\pi \in \Pi$, J_π^- is real-valued, or equivalently, if

$$\sup_{\pi\in\Pi} J_\pi^-(x) < \infty, \qquad \forall\, x \in S. \tag{2.1}$$

Note that the condition (2.1) requires in particular $g > -\infty$. The equivalence mentioned in the preceding definition can be verified directly. As a result of the condition (2.1), for an MDP in the (GC) class, we can define the *total cost of a policy* π in several equivalent ways:

$$\begin{aligned} J_\pi(x) &:= \limsup_{n\to\infty}\, \mathbb{E}_x^\pi \left\{\sum_{k=0}^{n} g(x_k,u_k)\right\} = \lim_{n\to\infty}\, \mathbb{E}_x^\pi \left\{\sum_{k=0}^{n} g(x_k,u_k)\right\} = \mathbb{E}_x^\pi \left\{\sum_{k=0}^{\infty} g(x_k,u_k)\right\} \\ &= \mathbb{E}_x^\pi \left\{\sum_{k=0}^{\infty} g_+(x_k,u_k)\right\} - \mathbb{E}_x^\pi \left\{\sum_{k=0}^{\infty} g_-(x_k,u_k)\right\} = J_\pi^+(x) - J_\pi^-(x). \end{aligned}$$

All these expressions are well-defined. In particular, under the condition (2.1), the sums of the one-stage costs, $\sum_{k=0}^{n} g(x_k,u_k)$ and $\sum_{k=0}^{\infty} g(x_k,u_k)$, along a trajectory $(x_0,u_0,x_1,u_1,\ldots)$, are well-defined with probability one under any policy. The situation $\infty-\infty$ does not occur in these summations (except on a negligible set of trajectories) or in any of the expectations above. The condition (2.1) is called the "general convergence condition" [15], from which comes the acronym (GC) for the class of problems satisfying this condition.

With the total cost of a policy defined as above, the *optimal cost function* J^* is given by

$$J^*(x) = \inf_{\pi\in\Pi} J_\pi(x) > -\infty, \qquad \forall\, x \in S,$$

where J^* does not take the value $-\infty$ by condition (2.1). A policy π is *optimal for state* x, if $J_\pi(x) = J^*(x)$. If this holds for all x, we say π is an *optimal policy*. Similarly, for $\epsilon > 0$, a policy π is *ϵ-optimal for state* x, if $J_\pi(x) \le J^*(x) + \epsilon$; and if this holds for all x, we say π is an *ϵ-optimal policy*.

Special Cases of the (GC) Model

As noted earlier, the (GC) class contains several special classes of total cost problems.[7] The case $g \ge 0$ is the positive costs model (P). The case $g \le 0$ is the negative costs model (N) with the additional condition that the optimal cost function is real-valued. The case where for each state x, there exists a control $u \in U(x)$ with $g(x,u) \le 0$, is the bounded negative costs model discussed in [28, Chap. 7.2] (where it is called positive bounded model in the reward maximization framework).

The (GC) model also covers a set of discounted problems with unbounded one-stage costs. Generally speaking, we can always reformulate discounted problems, with possibly transition-dependent discounting, as equivalent undiscounted total cost problems, using a standard procedure: Enlarge the state space to $S \cup \{\Delta\}$, where Δ represents an absorbing cost-free termination state, and then set the probability of transitioning to Δ from a state in S to be one minus the respective discount factor. For simplicity, in this paper we shall only consider a subset of the discounted problems that fall into the (GC) class. The discount factor in these problems is a given number, independent of the state transitions. We will refer to this class as (UD) in the rest of the paper.

[7] Some total cost problems are excluded from the (GC) model by the condition (2.1). As a simple example, if certain states can form a cycle in the state space under some policy, where some of the edges of the cycle have negative costs, then the negative edges will give rise to a total amount of $-\infty$ when the system goes through the cycle infinitely many times. This violates the condition (2.1), so even though the net cost may be nonnegative, such problems are not in the (GC) class. Shortest path or stochastic shortest path problems can be of this type. They can be analyzed as total cost problems under certain model conditions [3], when the average costs of all policies are nonnegative.

Definition 2.2 (The unbounded discounted model (UD))**.** Let $\beta \in [0,1)$. An MDP given above is in the class *(UD)* if, as in the definition of (GC), $\sup_{\pi\in\Pi} J_\pi^-(x) < \infty$ for all $x \in S$ and the objective is to minimize J_π over all policies $\pi \in \Pi$, but with the cost functions J_π^-, J_π defined by the β-discounted total costs of π as

$$J_\pi^-(x) = \mathbb{E}_x^\pi\left\{\sum_{k=0}^\infty \beta^k g_-(x_k,u_k)\right\}, \qquad J_\pi(x) = \mathbb{E}_x^\pi\left\{\sum_{k=0}^\infty \beta^k g(x_k,u_k)\right\}, \qquad \forall\, x\in S.$$

2.2 Some Optimality Properties of the (GC) Model

We discuss now some general properties of the optimal cost function and optimal policies for the (GC) model in the universal measurability framework. Since we focus exclusively on the (GC) class of total cost problems, in all the theorems given in this paper, we will not state the (GC) condition (2.1) explicitly.

Let $A(S)$ denote the set of lower semi-analytic functions $J : S \to [-\infty,+\infty]$, and let $\mathcal{M}(S)$ denote the set of universally measurable functions $J : S \to [-\infty,+\infty]$.[8] For the (GC) model, the dynamic programming operator $T : A(S) \to A(S)$ is defined by

$$T(J)(x) := \inf_{u\in U(x)}\left\{g(x,u) + \int_S J(x')\, q(dx' \mid x,u)\right\}, \qquad x\in S,\ \forall\, J\in A(S). \tag{2.2}$$

Associated with a stationary policy μ is the dynamic programming operator $T_\mu : \mathcal{M}(S) \to \mathcal{M}(S)$ defined by

$$T_\mu(J)(x) := \int_C \left\{g(x,u) + \int_S J(x')\, q(dx' \mid x,u)\right\}\, \mu(du \mid x), \qquad x\in S,\ \forall\, J\in\mathcal{M}(S).$$

In the above, the integral $\int_S J(x')\, q(dx' \mid x,u)$ appearing in T and T_μ is defined with respect to the completion of the Borel probability measure $q(\cdot \mid x,u)$ (see [2, p. 173] and see [10, Chap. 3.3] for the completion of measures), and the integral $\int_C\{\cdots\}\,\mu(du \mid x)$ in T_μ is defined likewise.

We clarify another detail in the definitions of T, T_μ given above, which concerns operations involving extended real-valued functions. In general, the integral $\int f d\rho$ for a Borel probability measure ρ and an extended real-valued, universally measurable function f is defined as $\int f d\rho = \int f_+ d\rho - \int f_- d\rho$, where f_+, f_- are the positive and negative parts of f, respectively. We use the convention $\infty - \infty = -\infty + \infty = \infty$ in this definition; we note, however, that such arithmetic operations will not be encountered for the (GC) model in this paper. In particular, under the (GC) condition (2.1), it can be verified that we do not encounter $\infty - \infty$ or $-\infty + \infty$ when computing $T(J^*)$, $T_\mu(J^*)$, or more generally, when computing $T(J)$ for any $J \in A(S)$ satisfying $J(x) \geq -\sup_{\pi\in\Pi} J_\pi^-(x)$, $x \in S$. At this point, however, there is still no need to introduce a general set of such functions J; such a set will be formally introduced later in Section 4.2.

We can state now some important properties of the (GC) model. For any policy π, $J_\pi \in \mathcal{M}(S)$ [2, Chap. 9.1], and for any stationary policy μ, J_μ satisfies the dynamic programming equation $J_\mu = T_\mu(J_\mu)$ (see Lemma A.1 in Appendix **??**). Furthermore, it can be shown that the (GC) model has the following optimality properties:

Theorem 2.1 (Properties of the optimal cost function)**.**

(a) The optimal cost function J^* is lower semi-analytic.

(b) The optimality equation $J^* = T(J^*)$ holds.

[8] A function $J : S \to [-\infty,+\infty]$ is universally measurable if it is measurable with respect to the universal σ-algebra on S. A lower semi-analytic function is universally measurable, so $A(S) \subset \mathcal{M}(S)$ [2, Chap. 7.7].

Theorem 2.2 (Properties of optimal policies)**.**

(a) For any $\epsilon > 0$, there exists an ϵ-optimal nonrandomized semi-Markov policy.

(b) If for each state an optimal policy exists, then there exists an optimal randomized semi-Markov policy.

(c) If $J^* \geq 0$, then there exists an ϵ-optimal nonrandomized Markov policy; if in addition, for each state an optimal policy exists, then there exists an optimal nonrandomized stationary policy. Furthermore, if $J^* \geq 0$, a stationary policy μ is optimal if and only if $J^* = T_\mu(J^*)$.

The preceding theorems largely follow from the same analyses given in the paper [33] and the book [2, Chap. 9.2-9.3] for positive or negative costs models. We will outline the proofs in Appendix B.

For problems with a discrete state space, proofs of Theorems 2.1(b), 2.2(a), as well as many other optimality results for (GC), can be found in the survey on total reward problems by Feinberg [15, Chap. 6].[9] (For the special positive and negative models, see also the books [2, Chap. 9] and [28, Chap. 7].)

We note that the existence of a *nonrandomized* ϵ-optimal semi-Markov policy, stated in Theorem 2.2(a), will be proved using a convergence property of value iteration in (GC) (cf. the discussion at the end of Section 3; see also the proof given in Appendix B.2.4). This is the only property, among the ones given by the preceding theorems, whose proof depends on the properties of value iteration. Our subsequent convergence analysis for value iteration will use the optimality equation $J^* = T(J^*)$ [Theorem 2.1(b)], so there is no circular reasoning involved.

The properties given in Theorem 2.2(c) are known to hold for the positive model (P) [2, Props. 9.12, 9.19], so Theorem 2.2(c) shows that the (GC) model is similar to (P) in some aspects when $J^* \geq 0$. In Section 4.1, we will give more results about value iteration in the case $J^* \geq 0$, which will show further connections between the two models. We also note that the last statement in Theorem 2.2(c) is implied by a more general statement in Schäl [29, Theorem 5.2.2], which does not require $J^* \geq 0$.

2.3 The Question about Value Iteration

In the rest of this paper we will focus on value iteration, i.e., the fixed point iteration $T^n(J)$ with the dynamic programming operator T, for some initial function $J \in A(S)$. The question of our interest is: for the (GC) model, under what conditions does $T^n(J)$ converge to J^*? The mode of convergence we consider is pointwise convergence. For a sequence of functions f_n, we shall consider the functions $\limsup_{n\to\infty} f_n$, $\liminf_{n\to\infty} f_n$, where the limsup and liminf are defined pointwise. When these two functions are equal, they define the pointwise limit function $\lim_{n\to\infty} f_n$. We write $f_n \to f$ or $\lim_{n\to\infty} f_n = f$ if f_n converges pointwise to a function f, and we write $f_n \uparrow f$ or $f_n \downarrow f$ if the convergence is monotonically from below or from above, respectively.

A natural function from which to start value iteration is the constant function zero, which we denote by $\mathbf{0}$. It can be a difficult starting point, however, as there exist many counterexamples with $T^n(\mathbf{0}) \not\to J^*$ (see e.g., [15]) and as our analysis in Section 4 will also show. An initial function greater than J^* can be easier for value iteration to converge. There is also evidence that in certain problems (such as deterministic shortest path problems), value iteration can converge significantly faster from above. We will start our analysis with such initial functions in the next section.

In this paper we focus on deriving convergence results for the (GC) model, without making further structural assumptions on the model. There are several classes of total cost MDP where additional structures can be exploited, and where the convergence of value iteration has been largely established. A major class of such problems is the one in which T can be shown to be a contraction

[9] The survey [15] focuses on GC MDP in which the state space S is countable, but the control space C can be an arbitrary measurable space.

operator on a certain complete metric space and the convergence of value iteration can be established with Banach's contraction principle. This class includes certain discounted problems as well as undiscounted problems with "transience structures" (see e.g., [20, Chaps. 8 and 9.6]). Semicontinuous models cover another broad class of problems for which value iteration has been relatively well studied (see e.g., the early work [29, Secs. 13-16] and the recent papers [16, Sec. 4], [17]), and for arbitrarily signed one-stage costs, the convergence of $T^n(\mathbf{0})$ has been shown under some additional conditions on the costs of policies (see e.g., [29, Condition (C)]). The results of this paper for the general (GC) model apply to semicontinuous models as well and can be useful as supplements to the existing results on value iteration developed by those early works just mentioned.

3 Convergence of Value Iteration from Above

In this section we prove a convergence theorem for value iteration starting from a function $J \geq J^*$. Let J^{*+} (resp. J^{*-}) be the optimal cost function if g_+ (resp. g_-) is the one-stage cost function:

$$J^{*+}(x) := \inf_{\pi\in\Pi} J_\pi^+(x), \qquad J^{*-}(x) := \inf_{\pi\in\Pi} J_\pi^-(x), \qquad x \in S.$$

Both functions are nonnegative, and J^{*+} is extended real-valued, whereas J^{*-} is real-valued under the model condition (2.1). Clearly $J^{*+} \geq J^{*+} - J^{*-} \geq J^*$.

Theorem 3.1 (Convergence of value iteration from above)**.** *For any function $J \in A(S)$ satisfying $J^* \leq J \leq c\big(J^{*+} + J^{*-}\big)$ for some $c \geq 1$, we have $T^n(J) \to J^*$.*

Theorem 3.1 extends a theorem of van der Wal, which states that $T^n(J) \to J^*$ if $J^* \leq J \leq J^{*+}$, in particular, $T^n(J^{*+}) \to J^*$ [37, Theorem 3.7] (see also [15, Theorem 6.8]). It can be seen from the definition of T and J^{*+} that when starting from J^{*+}, the convergence is monotonic: $T^n(J^{*+}) \downarrow J^*$. If the function $J^{*+} - J^{*-} \in A(S)$, then since $J^{*+} - J^{*-} \geq J^*$, it also satisfies the convergence condition and as can be easily verified, the convergence is also monotonic: $T^n\big(J^{*+} - J^{*-}\big) \downarrow J^*$.

Theorem 3.1 is also an extension of [41, Theorem 5.1] for the positive model (P), which states that for (P), $T^n(J) \to J^*$ if $J^* \leq J \leq cJ^*$ for some $c \geq 1$. This type of initial condition on value iteration resembles Whittle's bridging condition [40].

To prove Theorem 3.1, we will use arguments similar to those used in the proofs given in [41, Appendix E] for the model (P). The line of analysis was motivated by the work of Meyn [26, Chap. 9], who also systematically used the set of functions that are bounded by a multiple of the optimal cost function, to analyze average cost and total cost problems.

We will need the following two lemmas in the proof. Lemma 3.1 states some basic properties of the (GC) model, and Lemma 3.2 leads to the desired result. A subtlety here is that while the function J^{*+} is lower semi-analytic and therefore universally measurable [2, Chap. 7.7], the function J^{*-} is not necessarily universally measurable.[10] In the proof, we will need the outer integral of J^{*-}, which we introduce first.

For any nonnegative function J on S (possibly $J \notin \mathcal{M}(S)$, i.e., non-measurable), let $\widetilde{\mathbb{E}}_{n,x}^\pi\big\{J(x_n)\big\}$ denote the outer integral of J with respect to the marginal distribution of x_n induced by the policy π and the initial state $x_0 = x$; i.e.,

$$\widetilde{\mathbb{E}}_{n,x}^\pi\big\{J(x_n)\big\} := \inf\Big\{\, \mathbb{E}_x^\pi\big\{f(x_n)\big\} \,\Big|\, J \leq f,\ f \in \mathcal{M}(S)\Big\}. \tag{3.1}$$

By a property of the outer integral [2, Lemma A.2], if $J_1, J_2 \geq 0$ and $J_1 \in \mathcal{M}(S)$, then

$$\widetilde{\mathbb{E}}_{n,x}^\pi\big\{J_1(x_n) + J_2(x_n)\big\} = \mathbb{E}_x^\pi\big\{J_1(x_n)\big\} + \widetilde{\mathbb{E}}_{n,x}^\pi\big\{J_2(x_n)\big\}. \tag{3.2}$$

[10] Instead of being lower semi-analytic, the function $g_- = \max\{-g, 0\}$ is upper semi-analytic (i.e., $-g_-$ is lower semi-analytic). For this reason, J^{*-} may not be universally measurable (cf. [7, Example (48), p. 940]). If g is Borel measurable, then g_- would be lower semi-analytic and in turn, being the optimal cost function of the total cost problem with one-stage cost function g_-, J^{*-} would be lower semi-analytic and hence universally measurable.

Lemma 3.1 (Two basic properties)**.**

(a) For any given state x and policy π, there exists a Markov policy π' with

$$J^+_{\pi'}(x) = J^+_\pi(x), \qquad J^-_{\pi'}(x) = J^-_\pi(x), \qquad J_{\pi'}(x) = J_\pi(x).$$

(b) For any state x, policy π, and $n \geq 0$,

$$\mathbb{E}^\pi_x\left\{\sum_{k=n}^{\infty} g(x_k, u_k)\right\} \geq \mathbb{E}^\pi_x\left\{J^*(x_n)\right\}, \qquad \mathbb{E}^\pi_x\left\{\sum_{k=n}^{\infty} g_-(x_k, u_k)\right\} \geq \widetilde{\mathbb{E}}^\pi_{n,x}\left\{J^{*-}(x_n)\right\}.$$

We give the proof of the preceding lemma in Appendix A.

Lemma 3.2. *For a given $x \in S$, if π is a policy with $J_\pi(x) < +\infty$, then*

$$\lim_{n\to\infty} \mathbb{E}^\pi_x\big\{J^{*+}(x_n)\big\} = 0, \qquad \lim_{n\to\infty} \widetilde{\mathbb{E}}^\pi_{n,x}\big\{J^{*-}(x_n)\big\} = 0.$$

Proof. Since $J_\pi(x) = J^+_\pi(x) - J^-_\pi(x) < \infty$ by assumption and $J^-_\pi(x) < \infty$ by the model condition (2.1), we have $J^+_\pi(x) < \infty$. For any $n \geq 0$,

$$\begin{aligned} J^+_\pi(x) &= \mathbb{E}^\pi_x\left\{\sum_{k=0}^{n-1} g_+(x_k, u_k)\right\} + \mathbb{E}^\pi_x\left\{\sum_{k=n}^{\infty} g_+(x_k, u_k)\right\} \\ &\geq \mathbb{E}^\pi_x\left\{\sum_{k=0}^{n-1} g_+(x_k, u_k)\right\} + \mathbb{E}^\pi_x\big\{J^{*+}(x_n)\big\}. \end{aligned}$$

where the inequality follows from applying Lemma 3.1(b) to the total cost problem that has g_+ as the one-stage cost function. As $n \to \infty$, $\mathbb{E}^\pi_x\left\{\sum_{k=0}^{n-1} g_+(x_k, u_k)\right\} \uparrow J^+_\pi(x) < \infty$. So the preceding inequality implies that $\limsup_{n\to\infty} \mathbb{E}^\pi_x\big\{J^{*+}(x_n)\big\} = 0$ and hence $\lim_{n\to\infty} \mathbb{E}^\pi_x\big\{J^{*+}(x_n)\big\} = 0$. The second statement for J^{*-} is proved similarly. We start with the inequality

$$\begin{aligned} J^-_\pi(x) &= \mathbb{E}^\pi_x\left\{\sum_{k=0}^{n-1} g_-(x_k, u_k)\right\} + \mathbb{E}^\pi_x\left\{\sum_{k=n}^{\infty} g_-(x_k, u_k)\right\} \\ &\geq \mathbb{E}^\pi_x\left\{\sum_{k=0}^{n-1} g_-(x_k, u_k)\right\} + \widetilde{\mathbb{E}}^\pi_{n,x}\big\{J^{*-}(x_n)\big\}, \end{aligned}$$

which follows from Lemma 3.1(b). Since as $n \to \infty$, $\mathbb{E}^\pi_x\left\{\sum_{k=0}^{n-1} g_-(x_k, u_k)\right\} \uparrow J^-_\pi(x) < \infty$, we obtain from the preceding inequality $\lim_{n\to\infty} \widetilde{\mathbb{E}}^\pi_{n,x}\big\{J^{*-}(x_n)\big\} = 0$. □

Proof of Theorem 3.1. Let J be a function satisfying the condition of the theorem, i.e., $J \in A(S)$ and for some $c \geq 1$, $J^* \leq J \leq c(J^{*+} + J^{*-})$. Since $J \geq J^*$, $T^n(J) \geq T^n(J^*) = J^*$ for all n by the monotonicity of T. So to prove $T^n(J) \to J^*$, it is sufficient to consider an arbitrary state x with $J^*(x) < +\infty$ and show that $\limsup_{n\to\infty} T^n(J)(x) \leq J^*(x)$.

Let $\epsilon > 0$ and let $\pi = (\mu_0, \mu_1, \ldots)$ be an ϵ-optimal Markov policy for the state x; such a policy

exists by Lemma 3.1(a). For any $n \geq 1$, by the definition of T and the monotonicity of T, we have

$$\begin{aligned} T^n(J)(x) &\leq \big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_{n-1}}\big)(J)(x) \\ &= \mathbb{E}_x^\pi \left\{ \sum_{k=0}^{n-1} g(x_k, u_k) \right\} + \mathbb{E}_x^\pi \big\{ J(x_n) \big\} \\ &\leq \mathbb{E}_x^\pi \left\{ \sum_{k=0}^{n-1} g(x_k, u_k) \right\} + c\, \widetilde{\mathbb{E}}_{n,x}^\pi \big\{ J^{*+}(x_n) + J^{*-}(x_n) \big\}, \qquad (3.3) \\ &= \mathbb{E}_x^\pi \left\{ \sum_{k=0}^{n-1} g(x_k, u_k) \right\} + c\, \mathbb{E}_x^\pi \big\{ J^{*+}(x_n) \big\} + c\, \widetilde{\mathbb{E}}_{n,x}^\pi \big\{ J^{*-}(x_n) \big\}, \qquad (3.4) \end{aligned}$$

where we used the assumption $J \leq c\,(J^{*+} + J^{*-})$ and the definition (3.1) of outer integral in the inequality (3.3), and we used the fact $J^{*+} \in \mathcal{M}(S)$ and the equality (3.2) in Eq. (3.4). By the ϵ-optimality of π for x,

$$\lim_{n\to\infty} \mathbb{E}_x^\pi \left\{ \sum_{k=0}^{n-1} g(x_k, u_k) \right\} = J_\pi(x) \leq J^*(x) + \epsilon < \infty.$$

Applying Lemma 3.2 and using the fact $J_\pi(x) < \infty$, we also have $\lim_{n\to\infty} \mathbb{E}_x^\pi \big\{ J^{*+}(x_n) \big\} = 0$ and $\lim_{n\to\infty} \widetilde{\mathbb{E}}_{n,x}^\pi \big\{ J^{*-}(x_n) \big\} = 0$. Combining these relations with the inequality (3.4), we obtain

$$\limsup_{n\to\infty} T^n(J)(x) \leq J^*(x) + \epsilon.$$

Since ϵ is arbitrary, this gives $\limsup_{n\to\infty} T^n(J)(x) \leq J^*(x)$. Hence $T^n(J)(x) \to J^*(x)$ for all x, as discussed earlier. □

Remark 3.1. We have stated in Theorem 2.2(a) that for the (GC) model, an ϵ-optimal semi-Markov policy can be chosen to be nonrandomized. The construction of such a policy will use the fact that value iteration converges from above (for example, from J^{*+}). For details, see Appendix B.2.4. □

Let J^s be the minimal cost achievable by nonrandomized stationary policies:

$$J^s(x) := \inf_{\pi \in \Pi_s} J_\pi(x), \quad x \in S, \quad \text{where } \Pi_s = \big\{ \pi \in \Pi \mid \pi \text{ nonrandomized and statioanry} \big\}.$$

We say Π_s is *locally adequate* if $J^s(x) = J^*(x)$ for all $x \in S$ [31]. When the state space S is countable, it is known that $J^s = T(J^s)$ (Feinberg and Sonin [18, Theorem 2.2]; see also [15, Theorem 6.11]), so by Theorem 3.1 we obtain the following corollary about J^s.

Corollary 3.1. *Suppose S is countable. If Π_s is not locally adequate, i.e., $J^s \neq J^*$, then for any arbitrarily large $c > 0$, there exists some state x with $J^s(x) > c\big(J^{*+}(x) + J^{*-}(x)\big)$.*

For an example problem that illustrates Cor. 3.1, see e.g., [15, Example 6.5] (in which there exists a state x with a non-compact control set such that $J^s(x) = 1$ but $J^{*+}(x) + J^{*-}(x) = 0$). Let us also mention in passing that regarding the local adequacy of nonrandomized stationary policies, it is known that $J^s = J^*$ if S is countable and all the control sets $U(x)$ are finite (van der Wal [37, Theorem 2.22]), or if S is a Borel space and either $g \leq 0$ (Blackwell [5]) or certain continuity and compactness conditions are satisfied (Schäl [30]). For countable S, it is also known that $J^s = J^*$ under much weaker conditions (see Feinberg [15, Chaps. 6.7-6.9]).

The Special Case (UD)

For discounted problems in the (UD) class (cf. Def. 2.2), the dynamic programming operator T is given by Eq. (2.2), except that the integral in (2.2) is multiplied by the discount factor β:

$$T(J)(x) = \inf_{u \in U(x)} \left\{ g(x,u) + \beta \int_S J(x')\, q(dx' \mid x,u) \right\}, \qquad x \in S,\ \forall\, J \in A(S).$$

As noted by Whittle in the case of the positive costs model [40], for β-discounted problems, if $T^n(J) \to J^*$, then for any constant b, since $T^n(J+b) = T^n(J) + \beta^n b$, we also have $T^n(J+b) \to J^*$. Thus, the range of initial functions in Theorem 3.1 can be enlarged, leading to stronger conclusions:

Corollary 3.2. (UD) *For any function $J \in A(S)$ satisfying $J^* - b \le J \le c\,(J^{*+} + J^{*-}) + b$ for some scalars $c \ge 1$, $b \ge 0$, we have $T^n(J) \to J^*$. In particular, if J^* is bounded above, then $T^n(\mathbf{0}) \to J^*$.*

4 Convergence of Value Iteration starting from Zero

In this section we tackle the question of the convergence of $T^n(\mathbf{0})$ to J^*. The definition of J^* implies[11]

$$J_\infty := \limsup_{n\to\infty} T^n(\mathbf{0}) \le J^*, \tag{4.1}$$

so the convergence of $T^n(\mathbf{0})$ in question amounts to $\liminf_{n\to\infty} T^n(\mathbf{0}) \ge J^*$.

The behavior of $T^n(\mathbf{0})$ is fairly complex because the one-stage cost function can take both positive and negative values. It is known that for the (GC) model:

(i) $\{T^n(\mathbf{0})\}$ need not converge to J^*. It may not even have a pointwise limit. (See e.g., [15, Example 6.11] or the subsequent Example 5.1 in Section 5.)

(ii) Even if $J_\infty = \limsup_{n\to\infty} T^n(0)$ is a fixed point of T (and the limit of $T^n(\mathbf{0})$), it is still possible that $J_\infty \ne J^*$. This can happen even for finite state and control problems. (See e.g., [15, Example 6.10] or the subsequent Example 4.1 in Section 4.1.)

In what follows we first discuss two special cases, $J^* \le 0$ and $J^* \ge 0$, in which the behavior of $T^n(\mathbf{0})$ is shown to be similar to its behavior in the negative model and in the positive model, respectively. For the case $J^* \ge 0$, we will give several results about fixed point properties of T and convergence properties of transfinite value iteration. These results suggest that when $J^* \ge 0$, the (GC) model is similar to the positive costs model in many aspects. They lead us further to prove the convergence of non-transfinite value iteration for finite state or finite control problems, where value iteration employs the operator $\max\{\mathbf{0}, T(\cdot)\}$, which is derived from our analysis of transfinite value iteration.

We then focus on the general case where J^* may take both positive and negative values. We give sufficient conditions for the convergence of $T^n(\mathbf{0})$ and more generally, sufficient conditions for the convergence of $T^n(\mathbf{0})(x)$ to $J^*(x)$ for some states x. These conditions can be viewed as a generalization of the bridging condition of Whittle [40] from the positive cost model to the more general class of total cost problems considered here. We shall discuss this connection after presenting the convergence results. A special case of our results which relates to two existing sufficient conditions for value iteration will also be discussed then.

[11] For any policy π, $\limsup_{n\to\infty} T^n(\mathbf{0}) \le J_\pi$, from which the inequality (4.1) follows.

4.1 Special Cases where J^* is Nonpositive or Nonnegative

First, we observe that if $J^* \leq 0$, the problem is similar to the negative costs model (N), in the sense that $T^n(\mathbf{0}) \to J^*$. This follows from Theorem 3.1, but a direct proof is simpler: when $J^* \leq 0$, by the monotonicity of T, $T^n(\mathbf{0}) \geq T^n(J^*) = J^*$, whereas $\limsup_{n\to\infty} T^n(\mathbf{0}) \leq J^*$.

If $J^* \geq 0$, the (GC) model exhibits a number of properties analogous to those in the positive costs model (P), which will be the focus of our discussion below. As Theorem 2.2(c) already showed, in the (GC) model, when $J^* \geq 0$, we have the same structural properties of optimal policies as in the positive costs model. What we are going to discuss now are the fixed point properties of T and the convergence properties of value iteration starting from zero, when $J^* \geq 0$.

Let us first recall that in the positive costs model, if J is a nonnegative lower semi-analytic function on S and $J \geq T(J)$, then $J \geq J^*$, so J^* is the unique fixed point of T within the set $\{J \in A(S) \mid 0 \leq J \leq J^*\}$ [2, Prop. 9.10]. The (GC) model has a similar property:

Proposition 4.1. *If a function $J \in A(S)$ satisfies $J \geq 0$ and $J \geq T(J)$, then $J \geq J^*$.*

Proof. Under the (GC) condition (2.1), $T(\mathbf{0}) > -\infty$, so by the monotonicity of T and the assumption $J \geq 0$, $T(J)(x) \geq T(\mathbf{0})(x) > -\infty$ for all $x \in S$. It then follows from the selection theorem [2, Prop. 7.50] that given any $\epsilon > 0$, there exist stationary policies μ_k, $k \geq 0$, such that

$$T_{\mu_k}(J) \leq T(J) + 2^{-k-1}\epsilon \leq J + 2^{-k-1}\epsilon, \tag{4.2}$$

where the second inequality follows from the assumption $T(J) \leq J$. Consider the Markov policy $\pi = (\mu_0, \mu_1, \ldots)$. By Eq. (4.2) and the monotonicity of the mappings $T_{\mu_0}, T_{\mu_1}, \ldots$, we have

$$T_{\mu_0}(J) \leq J + 2^{-1}\epsilon, \qquad \big(T_{\mu_0} \circ T_{\mu_1}\big)(J) \leq J + (2^{-1} + 2^{-2}) \cdot \epsilon, \qquad \ldots$$

and consequently,

$$\big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_k}\big)(\mathbf{0}) \leq \big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_k}\big)(J) \leq J + \epsilon, \qquad \forall\, k \geq 1,$$

where we used the assumption $J \geq 0$ in the first inequality. As $k \to \infty$, the first expression in the preceding inequality converges to J_π in the (GC) model; therefore, $J^* \leq J_\pi \leq J + \epsilon$ and equivalently, $J^* \leq J$ since ϵ is arbitrary. □

The next two corollaries are direct consequences of the preceding proposition. The second corollary uses also the fact that $J_\infty = \limsup_{n\to\infty} T^n(\mathbf{0}) \leq J^*$.

Corollary 4.1. *If $J^* \geq 0$, then within the set $\{J \in A(S) \mid 0 \leq J \leq J^*\}$, J^* is the unique fixed point of T.*

Corollary 4.2. *If J_∞ is a fixed point of T and $J_\infty \geq 0$, then $J_\infty = J^*$. Hence, if J_∞ is a fixed point of T and yet $J_\infty \neq J^*$, then $J_\infty(x) < 0$ for some state x.*

For the discounted problems in the (UD) class (cf. Def. 2.2), by the same reasoning given at the end of Section 3, we can strengthen Cors. 4.1, 4.2 to obtain stronger conclusions. They are given by the next two corollaries, with the second one being a consequence of the first.

Corollary 4.3. (UD) *If $J^* \geq 0$, then J^* is the unique fixed point of T within the set*

$$\{J \in A(S) \mid -b \leq J \leq J^* + b, \text{ for some scalar } b \geq 0\}.$$

Proof. Let $\bar{J}$ be a fixed point of T in the set of functions stated by the corollary. Let β be the discount factor. We show that $\bar{J} = J^*$. Let $\hat{J} = \bar{J} + b \geq 0$. Then $T(\hat{J}) = T(\bar{J}) + \beta b = \bar{J} + \beta b \leq \hat{J}$, so by Prop. 4.1, $\hat{J} \geq J^*$. Consequently, $T^n(\hat{J}) \geq J^*$ for all n. We also have $T^n(\hat{J}) \to \bar{J}$ since $T^n(\hat{J}) = T^n(\bar{J}) + \beta^n b = \bar{J} + \beta^n b$. Therefore, $\bar{J} \geq J^*$. On the other hand, since $\bar{J} \leq J^* + b$, we have $\bar{J} = T^n(\bar{J}) \leq T^n(J^* + b) \to J^*$. Thus, $\bar{J} = J^*$. □

Corollary 4.4. (UD) *If $J^* \geq 0$ and if in addition, the function J_∞ is a fixed point of T and yet $J_\infty \neq J^*$, then $\inf_{x\in S} J_\infty(x) = -\infty$.*

Transfinite Value Iteration

Transfinite value iteration was introduced and analyzed by Maitra and Sudderth [24] for the positive costs model in the universal measurability framework. In an abstract dynamic programming context, it was also briefly discussed in [21, p. 463]. Transfinite value iteration is defined by transfinite recursion—recursion defined on a general well-ordered set; by comparison, value iteration is ordinarily defined by the classical recursion on the set of nonnegative integers. The need for studying transfinite value iteration comes from the convergence difficulty of ordinary value iteration.

The well-ordered set $(X, <)$ involved in transfinite value iteration is a set of ordinals, denoted $\{0, 1, \ldots, \omega_1\}$, with ω_1 being the first uncountable ordinal.[12] The transfinite value iteration for the positive costs model [24] can be described as follows: for ordinals $\xi \le \omega_1$, define recursively

$$T^0(J) = J; \qquad T^\xi(J) = T\Big(\sup_{\eta<\xi} T^\eta(J)\Big), \quad 0 < \xi < \omega_1; \qquad T^{\omega_1}(J) = \sup_{\xi<\omega_1} T^\xi(J). \tag{4.3}$$

It follows from the recursion principle on a well-ordered set [10, Theorem 1.3.2] and the properties of the universal measurability framework that the preceding recursion is well-defined, with $T^\xi(J) \in A(S)$ for all $\xi < \omega_1$.

We define transfinite value iteration similarly for the (GC) model. It will be more convenient for us to work with a mapping $\widetilde{T} : A(S) \to A(S)$, which is a modification of the dynamic programming operator T, and which has the desired increasing property $\widetilde{T}(J) \ge J$:

$$\widetilde{T}(J) := \max\big\{J, T(J)\big\}. \tag{4.4}$$

We use $\widetilde{T}$ to define transfinite value iteration as follows. For ordinals $\xi < \omega_1$, define recursively

$$\widetilde{T}^0(J) = J, \qquad \widetilde{T}^1(J) = \widetilde{T}(J),$$

and

$$\widetilde{T}^\xi(J) = \widetilde{T}\Big(\sup_{\eta<\xi} \widetilde{T}^\eta(J)\Big), \qquad 1 < \xi < \omega_1. \tag{4.5}$$

We then let

$$\widetilde{T}^{\omega_1}(J) = \sup_{\xi<\omega_1} \widetilde{T}^\xi(J). \tag{4.6}$$

Similar to the recursion (4.3), the preceding recursion is well-defined, with $\widetilde{T}^\xi(J) \in A(S)$ for all $\xi < \omega_1$, by the recursion principle on a well-ordered set [10, Theorem 1.3.2] and by the properties of the universal measurability framework.

Note that since $J^* = T(J^*)$, we have $\widetilde{T}(J^*) = J^*$ and hence $\widetilde{T}^{\omega_1}(J^*) = J^*$. Note also that by the increasing property $\widetilde{T}(J) \ge J$, the transfinite value iterates are nondecreasing:

$$\widetilde{T}^\xi(J) \ge \widetilde{T}^\eta(J) \quad \text{if} \ \ \xi > \eta.$$

Using the monotonicity of T, we can also write Eq. (4.5) equivalently as

$$\widetilde{T}^\xi(J) = \max\big\{J\,,\, T(\tilde{J}_\xi)\big\}, \qquad \text{where} \quad \tilde{J}_\xi = \sup_{\eta<\xi} \widetilde{T}^\eta(J). \tag{4.7}$$

For an integer $n \ge 0$, in general $\widetilde{T}^n(J) \ne T^n(J)$ (the ordinary value iteration). However, if J satisfies $T(J) \ge J$, then by the monotonicity of T, the two iterations become identical:

$$\widetilde{T}(J) = \max\big\{J, T(J)\big\} = T(J), \qquad \widetilde{T}^2(J) = \max\big\{\widetilde{T}(J), T(\widetilde{T}(J))\big\} = T^2(J), \quad \ldots,$$

[12] We give in Appendix C.1 an informal introduction of ordinals. As a full definition of ordinals is beyond the scope of this paper, we refer the readers to the books [10, Chaps. 1.3, A.3] and [22, p. 27-28] for the theory of ordinals, and their application in transfinite recursion and induction. See [10, Chaps. 1.3] also for the general principles of recursion and induction on well-ordered sets, which extend the principles of the classical recursion and induction on the set of nonnegative integers.

and so do $\widetilde{T}^{\xi}(J)$ and $T^{\xi}(J)$ for all ordinals $\xi \leq \omega_1$.

For the positive costs model (P), Maitra and Sudderth [24] showed that transfinite value iteration starting from zero converges to J^*, with $T^{\omega_1}(\mathbf{0}) = J^*$. Using their proof arguments together with Prop. 4.1, we can show that the (GC) model has a similar property:

Theorem 4.1 (Convergence of transfinite value iteration)**.** *If $J^* \geq 0$, then $\widetilde{T}^{\omega_1}(\mathbf{0}) = J^*$.*

Proof. Denote $J_{\omega_1} = \widetilde{T}^{\omega_1}(\mathbf{0})$. We use a critical lemma, Lemma C.1, given in Appendix C.2, which states that the function J_{ω_1} is lower semi-analytic and satisfies

$$J_{\omega_1} = \widetilde{T}(J_{\omega_1}) \geq T(J_{\omega_1}).$$

Combining this lemma with Prop. 4.1 and the assumption $J^* \geq 0$, we can obtain Theorem 4.1 immediately as follows. Since $J^* \geq 0$, using the definition of transfinite value iteration and the monotonicity of T, and using also the fact $J^* = T(J^*)$, we have

$$J_{\omega_1} = \widetilde{T}^{\omega_1}(\mathbf{0}) \leq \widetilde{T}^{\omega_1}(J^*) = J^*.$$

On the other hand, by the nondecreasing property of the transfinite value iterates, we have $J_{\omega_1} = \widetilde{T}^{\omega_1}(\mathbf{0}) \geq \widetilde{T}^{0}(\mathbf{0}) = \mathbf{0}$, and by Lemma C.1, we also have $J_{\omega_1} \geq T(J_{\omega_1})$ and J_{ω_1} is lower semi-analytic. So applying Prop. 4.1 with $J = J_{\omega_1}$, we have $J_{\omega_1} \geq J^*$, and therefore, we must have $J_{\omega_1} = J^*$. □

We mention that if the state and control spaces S, C are countable, then Theorem 4.1 holds with $\widetilde{T}^{\xi}(\mathbf{0}) = J^*$ for some countable ordinal $\xi < \omega_1$. This follows from the proof of [24, Theorem 5.1].[13]

If S is finite, then actually the sequence $\{\widetilde{T}^n(\mathbf{0})\}$ converges to J^* if J^* is nonnegative and real-valued. We prove this result below. It is similar to [41, Cor. 5.1], and the counterexample given immediately after [41, Cor. 5.1] in that paper for the (P) model is also applicable here and shows that the desired convergence need not hold if J^* can take the value $+\infty$. (For uncountable S, the proof given below is applicable if $\{\widetilde{T}^n(\mathbf{0})\}$ can be shown to converge uniformly.)

Proposition 4.2 (Convergence of value iteration for finite S)**.** *Suppose the state space S is finite. If $J^* \geq 0$ and J^* is real-valued, then $\widetilde{T}^n(\mathbf{0}) \to J^*$.*

Proof. By the definition of $\widetilde{T}$, the sequence $\{\widetilde{T}^n(\mathbf{0})\}$ is nondecreasing and converges monotonically to some limit $\bar{J} \geq 0$. Since $J^* \geq 0$, using the monotonicity of $\widetilde{T}$ and the fact $J^* = \widetilde{T}(J^*)$, we have $\bar{J} \leq J^*$, so $\bar{J} < +\infty$, in view of the assumption $J^* < +\infty$. Then, since S is finite and $\bar{J}$ is real-valued, the convergence $\widetilde{T}^n(\mathbf{0}) \to \bar{J}$ must be uniform; consequently, given any $\epsilon > 0$, $\bar{J} \leq \widetilde{T}^n(\mathbf{0}) + \epsilon$ for all n sufficiently large. Using this and the definition of $\widetilde{T}$, we obtain for any given $\epsilon > 0$ that for all n sufficiently large,

$$\widetilde{T}\big(\bar{J}\big) \leq \widetilde{T}\big(\widetilde{T}^n(\mathbf{0}) + \epsilon\big) \leq \widetilde{T}^{n+1}(\mathbf{0}) + \epsilon.$$

This implies $\widetilde{T}(\bar{J}) \leq \bar{J}$; hence by Prop. 4.1, $\bar{J} \geq J^*$. But $\bar{J} \leq J^*$ as discussed earlier, so $\bar{J} = J^*$. □

We have used $\widetilde{T}$ instead of T to define transfinite value iteration for the (GC) model. It is not hard to verify that the transfinite value iteration defined by Eq. (4.3) is essentially equivalent to the iteration (4.5)-(4.6), and therefore, by Theorem 4.1, $T^{\omega_1}(\mathbf{0}) = J^*$ if $J^* \geq 0$. However, if instead we define transfinite value iteration more similarly to ordinal value iteration as

$$\hat{J}_1 = T(\mathbf{0}), \qquad \hat{J}_\xi = T\Big(\sup_{1 \leq \eta < \xi} \hat{J}_\eta\Big), \quad 1 < \xi < \omega_1, \qquad \hat{J}_{\omega_1} = \sup_{1 \leq \xi < \omega_1} \hat{J}_\xi,$$

[13] In fact, as can be verified from its proof, the conclusion of [24, Theorem 5.1] also holds in our case when S or C is uncountably infinite. Specifically, we have that $\widetilde{T}^{\xi}(\mathbf{0}) = J^*$ for some $\xi < \omega_1$ if $J^* \geq 0$ and if there exists a probability measure $\rho \in \mathcal{P}(S)$ such that for every $(x, u) \in S \times C$, $q(dx' \mid x, u)$ is absolutely continuous with respect to ρ.

would we still have $\hat{J}_{\omega_1} = J^*$ if $J^* \geq 0$? The next example shows that it is not true in general, even if S, C are finite.[14] This example is a slight modification of [37, Example 3.2] (cf. [15, Example 6.10]).

Example 4.1. Let $S = \{0, 1, 2\}$. State 0 is cost-free and absorbing. From state 1, the system moves to state 0 deterministically, with cost 1. State 2 has two controls: under one control the system remains at state 2 with no cost, and under the other the system moves to state 1 with cost -1. The optimal costs at these states are $J^*(0) = 0, J^*(1) = 1, J^*(2) = 0$, so $J^* \geq 0$. However, for state 2, $\hat{J}_{\omega_1}(2) = -1$, since for all $1 \leq \xi < \omega_1$, $\hat{J}_\xi(2) = -1$. Indeed, the function $\hat{J}_1 = T(\mathbf{0})$, which has values $0, 1, -1$ for states $0, 1, 2$, respectively, is a fixed point of T. But it is easy to see that $\widetilde{T}(\mathbf{0}) = J^*$. □

The next corollary follows from Theorem 4.1. (Note that $J_\infty \in A(S)$ by [2, Lemma 7.30(2)].)

Corollary 4.5. *If $J_\infty \geq 0$, then $\widetilde{T}^{\omega_1}(J_\infty) = J^*$.*

Proof. Since $J^* \geq J_\infty \geq 0$, using Theorem 4.1 and the monotonicity of transfinite value iteration, we have that $J^* = \widetilde{T}^{\omega_1}(\mathbf{0}) \leq \widetilde{T}^{\omega_1}(J_\infty) \leq J^*$. This shows $\widetilde{T}^{\omega_1}(J_\infty) = J^*$. □

Finally, for the case $J^* \geq 0$, we prove the convergence of the sequence $\{\widetilde{T}^n(\mathbf{0})\}$ when the control sets are finite. This result can be compared with the convergence of $T^n(\mathbf{0})$ in positive costs problems with finite controls (cf. [1], [2, Prop. 9.18(P)]).

Proposition 4.3 (Convergence of value iteration in the case of finite controls)**.** *Suppose that the control set $U(x)$ is finite for each $x \in S$. Then $\widetilde{T}^n(\mathbf{0}) \to J^*$ if $J^* \geq 0$.*

Proof. Denote $J_n = \widetilde{T}^n(\mathbf{0})$ with $J_0 = \mathbf{0}$. Similarly to the proof of Prop. 4.2, we have that $J_n \geq 0$ and $J_n \uparrow \bar{J}$ for some nonnegative function $\bar{J} \in A(S)$ with $\bar{J} \leq J^*$. Consider an arbitrary state x. Since $U(x)$ is finite, there exist a control $\bar{u} \in U(x)$ and a subsequence $\{n_k\}$ of positive integers such that for all $k \geq 1$, $\bar{u}$ is the control that attains the minimum in $T\big(J_{n_k}\big)(x)$, i.e.,

$$J_{n_k+1}(x) = \widetilde{T}\big(J_{n_k}\big)(x) = \max\big\{J_{n_k}(x),\, T\big(J_{n_k}\big)(x)\big\} = \max\big\{J_{n_k}(x),\, H\big(x, \bar{u}, J_{n_k}\big)\big\}, \tag{4.8}$$

where

$$H\big(x, \bar{u}, J_{n_k}\big) := g(x, \bar{u}) + \int_S J_{n_k}(x')\, q(dx' \mid x, \bar{u}) = \inf_{u \in U(x)} \left\{ g(x, u) + \int_S J_{n_k}(x')\, q(dx' \mid x, u) \right\}.$$

Since $J_{n_k} \geq 0$ and $J_{n_k} \uparrow \bar{J}$, we have $\lim_{k\to\infty} \int_S J_{n_k}(x')\, q(dx' \mid x, \bar{u}) = \int_S \bar{J}(x')\, q(dx' \mid x, \bar{u})$ by the monotone convergence theorem. Hence, from Eq. (4.8), by letting k go to $+\infty$ and then applying the inequality $g(x, \bar{u}) + \int_S \bar{J}(x')\, q(dx' \mid x, \bar{u}) \geq T(\bar{J})(x)$, we have

$$\bar{J}(x) \geq \max\big\{\bar{J}(x),\, T\big(\bar{J}\big)(x)\big\} \quad \Longrightarrow \quad \bar{J}(x) \geq T\big(\bar{J}\big)(x).$$

Since x is arbitrary, this proves $\bar{J} \geq T(\bar{J})$, which together with the fact $\bar{J} \geq 0$ implies $\bar{J} \geq J^*$ by Prop. 4.1. But $\bar{J} \leq J^*$ as we showed earlier, so $\bar{J} = J^*$. □

Proposition 4.3 gives us a practical way of computing J^* for finite control problems when we can infer $J^* \geq 0$ from the problem data relatively easily without intensive computation. We also note that the iteration $J_{n+1} = \widetilde{T}(J_n)$ with $J_0 = \mathbf{0}$ can be equivalently written as the slightly simpler iteration $J_{n+1} = \max\{\mathbf{0}, T(J_n)\}$, as mentioned earlier in Eq. (4.7) and can be verified by induction.

[14] Note, however, that finite state and control MDP under the total cost criterion can be satisfactorily solved using a policy iteration algorithm based on the sensitive optimality concept [28, Chap. 10.4].

4.2 The General Case

We now consider the convergence of $T^n(\mathbf{0})$ in the general case where J^* can take both positive and negative values. We will formulate sufficient convergence conditions, using some ideas from Whittle's bridging condition for the positive costs model (Whittle [40], Hartley [19]). Our conditions can characterize convergence of $T^n(\mathbf{0})$ on parts of the state space, so they can be useful even if on the full state space, $T^n(\mathbf{0}) \not\to J^*$. These conditions can be viewed as generalizations of the bridging condition from the positive model to the more general (GC) model, and we will discuss this connection after giving our convergence results (see Remark 4.1 in Section 4.2.3).

We start by introducing a few mappings, function spaces and inequalities, which we will need in analyzing the convergence of $\{T^n(\mathbf{0})\}$ and bounding these iterates from below.

4.2.1 Preliminaries

Define two dynamic programming operators $T_+ : A(S) \to A(S)$, $T_- : A(S) \to A(S)$, corresponding to the positive and negative parts of g, respectively, as follows: for any $J \in A(S)$,

$$T_+(J)(x) := \inf_{u \in U(x)} \left\{ g_+(x,u) + \int_S J(x')\, q(dx' \mid x,u) \right\}, \qquad x \in S,$$
$$T_-(J)(x) := \inf_{u \in U(x)} \left\{ -g_-(x,u) + \int_S J(x')\, q(dx' \mid x,u) \right\}, \qquad x \in S.$$

Also define $T_0 : A(S) \to A(S)$ by

$$T_0(J)(x) := \inf_{u \in U(x)} \int_S J(x')\, q(dx' \mid x,u), \qquad x \in S,\ \forall\, J \in A(S).$$

We may treat $\infty - \infty = \infty$ in the above definitions. However, for the inequalities we will need, we will focus on a smaller set of functions, $A_0(S) \subset A(S)$, so that $\infty - \infty$ does not occur when applying any of the above mappings to a function in the set. We define this set as follows.

Let $A_0(S)$ be a subset of lower semi-analytic functions given by

$$A_0(S) := \left\{ J \in A(S) \;\middle|\; J > -\infty,\ \inf_{n \geq 1} T_0^n\big(\min\{J, 0\}\big) > -\infty \right\}, \tag{4.9}$$

where the infimum over n is taken pointwise. The set $A_0(S)$ has the desired property that if $J_1, J_2 \in A_0(S)$, then in the calculation of $J_1 + J_2$ or the integral $\int_S \big(J_1(x') + J_2(x')\big)\, q(dx' \mid x,u)$, we never encounter $\infty - \infty$ or $-\infty + \infty$, and moreover, with $J = J_1 + J_2$, the linearity property of expectation can be taken for granted:[15]

$$\int_S J(x')\, q(dx' \mid x,u) = \int_S J_1(x')\, q(dx' \mid x,u) + \int_S J_2(x')\, q(dx' \mid x,u).$$

The set $A_0(S)$ contains all real-valued constant functions and all nonnegative lower semi-analytic functions. Under the (GC) condition (2.1), $A_0(S)$ also contains the following set of functions:

$$\{J \in A(S) \mid J \geq f\}, \qquad \text{where } f(x) = -\sup_{\pi \in \Pi} J_\pi^-(x), \quad x \in S, \tag{4.10}$$

[15] This is not the case if J_1 or J_2 is not in the set $A_0(S)$. As an example, suppose $J_1 = f$ and $J_2 = -f$, where $\int_S f(x')\, q(dx' \mid x,u) = \infty$, so that $J_2 \notin A_0(S)$. Then $0 = f - f$, but

$$0 = \int_S 0 \cdot q(dx' \mid x,u) \neq \int_S f(x')\, q(dx' \mid x,u) + \int_S -f(x')\, q(dx' \mid x,u) = +\infty - \infty = +\infty.$$

The definition of $A_0(S)$ rules out such cases.

because the nonpositive function $f \in A(S)$ satisfies that $T_0(f) \geq T_-(f) = f > -\infty$ under the condition (2.1).[16] In particular, $J^*, J_\infty \in A_0(S)$.

It can be verified directly that the set $A_0(S)$ is closed under addition and positive scaling: if $J_1, J_2 \in A_0(S)$, then $c_1 J_1 + c_2 J_2 \in A_0(S)$ for any scalars $c_1, c_2 \geq 0$. It can also be verified that the set $A_0(S)$ is closed under the mappings T and T_+, T_-, T_0 defined above, under the (GC) condition (2.1).[17]

For functions in $A_0(S)$, the inequalities given in the following lemma can be verified by a direct calculation, using the definitions of the mappings T, T_+, T_-, T_0 and the monotonicity of T. The inequality (4.12) below shows a concavity property of T.

Lemma 4.1. *Let $J_1, J_2 \in A_0(S)$. Then*

$$T^n(J_1 + J_2) \geq T_+^n(J_1) + T_-^n(J_2), \qquad T^n(J_1 + J_2) \geq T^n(J_1) + T_0^n(J_2), \qquad n \geq 1, \tag{4.11}$$

and for any $\alpha \in [0,1]$,

$$T^n\big(\alpha J_1 + (1-\alpha)J_2\big) \geq \alpha\, T^n(J_1) + (1-\alpha)\, T^n(J_2), \qquad n \geq 1. \tag{4.12}$$

4.2.2 Convergence Results

We give two convergence theorems below. They concern the convergence of $T^n(\mathbf{0})(x)$ to $J^*(x)$ for *some* states x. We will discuss their connection to Whittle's bridging condition and also derive several corollaries in the subsequent Section 4.2.3, and we will then use examples to illustrate these theorems in Section 5.

Theorem 4.2 (Partial convergence of value iteration)**.** *Suppose that*

$$T^{\bar n}(\mathbf{0}) \geq \alpha J^* + \phi \quad \text{for some } \phi \in A_0(S),\ \bar n \geq 0,\ \alpha \in (0,1]. \tag{4.13}$$

Then

$$\lim_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x), \quad \forall\, x \in S_{0^+} := \Big\{ x \in S \,\Big|\, \limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} \leq 0 \Big\}.$$

In addition, if $\bar n = 0$, then

$$\limsup_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x), \quad \forall\, x \in S_0 := \Big\{ x \in S \,\Big|\, \liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} \leq 0 \Big\};$$

if $\bar n \neq 0$, then for every $x \in S_0$ such that $T^n(\mathbf{0})(x)$ converges, $\lim_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x)$.

Theorem 4.3 (Partial convergence of value iteration)**.** *Suppose that*

$$J_\infty = T(J_\infty) \quad \text{and} \quad J_\infty \geq \alpha J^* + \phi \quad \text{for some } \phi \in A_0(S),\ \alpha \in (0,1]. \tag{4.14}$$

Then $J_\infty(x) = J^*(x)$ for all $x \in S_{0^+}$ and for all $x \in S_0$ such that $T^n(\mathbf{0})(x)$ converges, where the sets S_{0^+}, S_0 are as given in Theorem 4.2.

[16] We have $f \in A(S)$ and $T_-(f) = f$ because f is the optimal cost function of the nonpositive total cost problem with one-stage cost function $-g_-$.

[17] Denote $\lceil J \rceil^0 = \min\{J, 0\}$ for any function J. The set $A_0(S)$ consists of those functions $J \in A(S)$ such that all the functions $T_0^n(\lceil J \rceil^0), n = 0, 1, \ldots$, lie above some function f_J which does not take the value $-\infty$. This set $A_0(S)$ is by definition closed under the mapping T_0, and hence, since $T_+(J) \geq T_0(J)$ for any function J, $A_0(S)$ is also closed under T_+. Similarly, since $T(J) \geq T_-(J)$, if $A_0(S)$ is closed under T_-, then it must also be closed under T. To show that $A_0(S)$ is closed under T_-, we use the (GC) condition (2.1) as follows. Let $f \leq 0$ be the function given in Eq. (4.10). We have $f > -\infty$ under the condition (2.1) and $T_0(f) \geq T_-(f) = f$. Then for $J \in A_0(S)$, by a direct calculation, we have

$$T_-(J) \geq f + T_0\big(\lceil J \rceil^0\big) > -\infty, \qquad T_0^n\big(\lceil T_-(J) \rceil^0\big) \geq f + T_0^{n+1}\big(\lceil J \rceil^0\big) \geq f + f_J > -\infty, \quad \forall\, n \geq 1,$$

so $T_-(J) \in A_0(S)$. This shows that the set $A_0(S)$ is closed under the mappings T_- and T.

In the above theorems we have expressed the conditions in terms of the optimal cost function J^* itself. When applying the theorems, we can always relax the conditions and replace J^* with the cost function of some policy, so that the conditions can be verified in practice (cf. the discussion in Section 4.2.3). The function ϕ appearing in the conditions above reflects the difference between the optimal cost function and the iterate $T^{\bar{n}}(\mathbf{0})$ or the limit function J_∞ from value iteration. Both theorems use ϕ to infer the convergence of value iteration, in a manner that resembles some other sufficient conditions in the literature, which are based on the "tail behavior" of policies (cf. the subsequent Prop. 4.4 and the related discussion in Section 4.2.3).

The proofs of Theorems 4.2 and 4.3 are similar. First, note that for each $x \in S$, the definition of T_0 implies $T_0^n(\phi)(x) = -\sup_{\pi\in\Pi} E_x^\pi\{-\phi(x_n)\}$, so the sets S_{0^+}, S_0 are equivalently given by

$$S_{0^+} = \Big\{x \in S \,\Big|\, \liminf_{n\to\infty} T_0^n(\phi)(x) \geq 0\Big\}, \qquad S_0 = \Big\{x \in S \,\Big|\, \limsup_{n\to\infty} T_0^n(\phi)(x) \geq 0\Big\}. \tag{4.15}$$

Proof of Theorem 4.2. For any $m \geq 0$, $T^m\big(T^{\bar{n}}(\mathbf{0})\big) \geq T^m(\alpha J^* + \phi)$ by the assumption (4.13) and the monotonicity of T. Since $\phi \in A_0(S)$ by assumption and $\alpha J^* \in A_0(S)$ under the (GC) condition (2.1), applying the second inequality in (4.11) with $J_1 = \alpha J^*, J_2 = \phi$, we have

$$\begin{aligned} T^{m+\bar{n}}(\mathbf{0}) &\geq T^m(\alpha J^*) + T_0^m(\phi) \\ &\geq \alpha J^* + (1-\alpha)\, T^m(\mathbf{0}) + T_0^m(\phi), \end{aligned} \tag{4.16}$$

where in the last inequality we applied the inequality (4.12) with $J_1 = J^*, J_2 = \mathbf{0}$ and used the fact $T^m(J^*) = J^*$.

Consider the inequality (4.16) for each $x \in S$ separately. Note that under the condition (2.1), we have $J^*(x) > -\infty$ and $\inf_{m\geq 1} T^m(\mathbf{0})(x) \geq f(x) > -\infty$ where the function f is as given in Eq. (4.10). We also have, by the assumption $\phi \in A_0(S)$, that $\inf_{m\geq 1} T_0^m(\phi)(x) > -\infty$. Hence, from the inequality (4.16) for each state x, by taking liminf of both sides, we obtain that for each $x \in S$,

$$\liminf_{n\to\infty} T^n(\mathbf{0})(x) \geq \alpha J^*(x) + (1-\alpha) \liminf_{n\to\infty} T^n(\mathbf{0})(x) + \liminf_{n\to\infty} T_0^n(\phi)(x).$$

Since $\liminf_{n\to\infty} T^n(\mathbf{0})(x) > -\infty$ and $\alpha \in (0, 1]$, this implies

$$\liminf_{n\to\infty} T^n(\mathbf{0})(x) \geq J^*(x) + \tfrac{1}{\alpha} \liminf_{n\to\infty} T_0^n(\phi)(x). \tag{4.17}$$

For $x \in S_{0^+}$, Eqs. (4.15), (4.17) together imply

$$\liminf_{n\to\infty} T^n(\mathbf{0})(x) \geq J^*(x), \quad \forall\, x \in S_{0^+};$$

but $\limsup_{n\to\infty} T^n(\mathbf{0})(x) \leq J^*(x)$ always. Hence $\lim_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x)$ for all $x \in S_{0^+}$.

The proof of the statements for $x \in S_0$ is similar. For each $x \in S$, take a subsequence $\{n_k\}$ of integers such that $\lim_{k\to\infty} T_0^{n_k}(\phi)(x) = \limsup_{n\to\infty} T_0^n(\phi)(x)$. We have, by inequality (4.16),

$$T^{n_k+\bar{n}}(\mathbf{0})(x) \geq \alpha J^*(x) + (1-\alpha)\, T^{n_k}(\mathbf{0})(x) + T_0^{n_k}(\phi)(x), \qquad k \geq 1.$$

As we showed earlier, $J^*(x) > -\infty$, $\inf_{m\geq 1} T^m(\mathbf{0})(x) > -\infty$, and $\inf_{m\geq 1} T_0^m(\phi)(x) > -\infty$. So, by taking limsup of both sides of the preceding inequality, we obtain

$$\limsup_{k\to\infty} T^{n_k+\bar{n}}(\mathbf{0})(x) \geq \alpha J^*(x) + (1-\alpha) \limsup_{k\to\infty} T^{n_k}(\mathbf{0})(x) + \limsup_{n\to\infty} T_0^n(\phi)(x). \tag{4.18}$$

If $\bar{n} = 0$ and $x \in S_0$, then Eqs. (4.15), (4.18) together imply

$$\limsup_{k\to\infty} T^{n_k}(\mathbf{0})(x) \geq \alpha J^*(x) + (1-\alpha) \limsup_{k\to\infty} T^{n_k}(\mathbf{0})(x),$$

and since $\limsup_{k\to\infty} T^{n_k}(\mathbf{0})(x) > -\infty$ and $\alpha \in (0,1]$, we then have

$$\limsup_{k\to\infty} T^{n_k}(\mathbf{0})(x) \geq J^*(x) \quad \Longrightarrow \quad \limsup_{n\to\infty} T^n(\mathbf{0})(x) \geq J^*(x),$$

which is equivalent to $\limsup_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x)$ [cf. Eq. (4.1)]. If $\bar{n} \neq 0$ but $x \in S_0$ and $T^n(\mathbf{0})(x)$ converges, then the inequality (4.18) together with Eq. (4.15) implies that

$$\lim_{n\to\infty} T^n(\mathbf{0})(x) \geq \alpha J^*(x) + (1-\alpha) \lim_{n\to\infty} T^n(\mathbf{0})(x),$$

which in turn implies $\lim_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x)$, similar to the preceding proof. □

Proof of Theorem 4.3. Recall that $J_\infty \leq J^*$, so to show $J_\infty(x) = J^*(x)$ for some state x, it is sufficient to show $J_\infty(x) \geq J^*(x)$.

Using the condition (4.14): $J_\infty \geq \alpha J^* + \phi$ where $\phi \in A_0(S)$, and using also Lemma 4.1, we have, similar to the proof of Theorem 4.2, that for any $m \geq 0$,

$$J_\infty = T^m(J_\infty) \geq T^m(\alpha J^* + \phi) \geq \alpha J^* + (1-\alpha)\, T^m(\mathbf{0}) + T_0^m(\phi). \tag{4.19}$$

Consider this inequality for each $x \in S$ separately. For each x, take a subsequence $\{n_k\}$ of integers such that $\lim_{k\to\infty} T^{n_k}(\mathbf{0})(x) = \limsup_{n\to\infty} T^n(\mathbf{0})(x)$. By the inequality (4.19), we have

$$J_\infty(x) \geq \alpha J^*(x) + (1-\alpha)\, T^{n_k}(\mathbf{0})(x) + T_0^{n_k}(\phi)(x), \qquad \forall\, k \geq 1. \tag{4.20}$$

As discussed in the proof of Theorem 4.2, we have $J^*(x) > -\infty$, $\inf_{m\geq 1} T^m(\mathbf{0})(x) > -\infty$, and $\inf_{m\geq 1} T_0^m(\phi)(x) > -\infty$. Consequently, by letting k go to ∞ in the preceding inequality, we obtain

$$\begin{aligned} J_\infty(x) &\geq \alpha J^*(x) + (1-\alpha) \limsup_{n\to\infty} T^n(\mathbf{0})(x) + \liminf_{n\to\infty} T_0^n(\phi)(x) \\ &\geq \alpha J^*(x) + (1-\alpha) J_\infty(x) + \liminf_{n\to\infty} T_0^n(\phi)(x), \end{aligned} \tag{4.21}$$

and since $J_\infty(x) > -\infty$ and $\alpha \in (0,1]$, this implies

$$J_\infty(x) \geq J^*(x) + \tfrac{1}{\alpha} \liminf_{n\to\infty} T_0^n(\phi)(x).$$

If $x \in S_{0^+}$, the preceding inequality together with Eq. (4.15) yields $J_\infty(x) \geq J^*(x)$, which implies $J_\infty(x) = J^*(x)$, as discussed earlier. This proves the statement for $x \in S_{0^+}$.

From the inequality (4.19), we also have that for any $x \in S$,

$$J_\infty(x) \geq \alpha J^*(x) + (1-\alpha) \liminf_{n\to\infty} T^n(\mathbf{0})(x) + \limsup_{n\to\infty} T_0^n(\phi)(x). \tag{4.22}$$

This inequality is obtained similarly to the inequality (4.21), by starting from the inequality (4.20) but with the subsequence $\{n_k\}$ being such that $\lim_{k\to\infty} T_0^{n_k}(\phi)(x) = \limsup_{n\to\infty} T_0^n(\phi)(x)$. If $x \in S_0$ and $T^n(\mathbf{0})(x)$ converges so that $J_\infty(x) = \lim_{n\to\infty} T^n(\mathbf{0})(x)$, then the inequality (4.22) together with Eq. (4.15) implies $J_\infty(x) \geq \alpha J^*(x) + (1-\alpha) J_\infty(x)$, which is equivalent to $J_\infty(x) \geq J^*(x)$ (since $J_\infty(x) > -\infty$ and $\alpha \in (0,1]$), and this is in turn equivalent to $J_\infty(x) = J^*(x)$, as discussed earlier. This completes the proof. □

4.2.3 Implications and Discussion

We start with two direct consequences of Theorem 4.2. The first corollary simply relaxes the condition of Theorem 4.2 using the fact $J^* \leq J_\pi$ for any policy π; the relaxed condition is easier to verify in practice.

Corollary 4.6. *Suppose there exists a policy π satisfying $T^{\bar{n}}(\mathbf{0}) \geq \alpha J_\pi + \phi$ for some $\phi \in A_0(S)$, $\bar{n} \geq 0$, and $\alpha \in (0,1]$. Then the conclusions of Theorem 4.2 hold.*

To prepare for a discussion on the relation of our results with the bridging condition, let us state another evident implication of Theorem 4.2 by letting $\phi(\cdot) \equiv 0$ in the theorem.

Corollary 4.7. *We have $T^n(\mathbf{0}) \to J^*$ if there exist $\bar{n} \geq 0$ and $\alpha \in (0,1]$ such that*

$$T^{\bar{n}}(\mathbf{0}) \geq \alpha J^* \quad \text{or} \quad T^{\bar{n}}(\mathbf{0}) \geq \alpha J_\pi \quad \text{for some policy } \pi. \tag{4.23}$$

Remark 4.1. We can compare the condition of Cor. 4.6 with the bridging condition for the positive model (P) [40], which can be stated in our notation as $T^{\bar{n}}(\mathbf{0}) \geq \alpha J_\pi$ for some policy π, i.e., the condition (4.23) in Cor. 4.7. According to our understanding of [40], the name "bridging" reflects the fact that the condition relates the finite-horizon cost $T^{\bar{n}}(\mathbf{0})$ to the infinite-horizon cost J_π. In this sense, the conditions of our theorems are also of the "bridging-type," and in our theorems, the differences between the finite-horizon and infinite-horizon costs are reflected in the function ϕ and used to infer the convergence of value iteration at various states. For the case $g \geq 0$, Whittle [40] and then Hartley [19] gave two different proofs, establishing that under the bridging condition, $T^n(J) \to J^*$ for all functions J with $0 \leq J \leq cJ^*$ for some real c. Our proofs of Theorems 4.2, 4.3 indeed started out similarly to Hartley's.

For the (GC) model, we see that under the condition (4.23), the desired convergence $T^n(\mathbf{0}) \to J^*$ holds. Thus it may seem that we could have taken the original bridging condition (4.23), which is much simpler than the conditions of Theorem 4.2 or Cor. 4.6, for the more general total cost problems in (GC). But the fact is that condition (4.23) can be not only restrictive but also unnatural for the (GC) model, where J^* can take negative values. In particular, for those states x with $J^*(x) \leq 0$, $\alpha J^*(x) \geq J^*(x)$ for all $\alpha \in (0,1]$, whereas it is possible that $T^n(\mathbf{0})(x) < J^*(x)$ for all n and then the inequalities in (4.23) cannot be satisfied. This is the reason for introducing the function ϕ in formulating the conditions of Theorems 4.2, 4.3.

If the conditions of Theorem 4.2 are met with $S_{0^+} = S$, we have $T^n(\mathbf{0}) \to J^*$, and this combined with Theorem 3.1 then implies, by the monotonicity of T, that $T^n(J) \to J^*$ for all lower semi-analytic functions J such that $0 \leq J \leq c(J^{*+} + J^{*-})$ for some real c. This set of J then resembles the set of initial functions in the conclusion of [40] for the model (P) mentioned earlier, although here J^* need not lie in this set because it can take negative values. □

We now discuss another special case of Theorem 4.2 and relate it to two existing sufficient conditions for the convergence of value iteration. Take $\bar{n} = 0$ in Theorem 4.2 with $\phi = -J^*$, $\alpha = 1$. We deduce that if J^* is real-valued and $-J^* \in A_0(S)$, then the conclusions given in the subsequent Prop. 4.4 must hold. However, in this case, these conditions on J^* turn out to be unnecessary (they are artifacts of the limitation of the proof technique we used), as a short direct proof given below shows. This result, Prop. 4.4, relates to a theorem of van Hee et al. [38, Theorem 3.5]: Indeed the first half of Prop. 4.4 is a "pointwise version" of [38, Theorem 3.5], and the proof argument for Prop. 4.4, included below for completeness, is essentially from [38].

Proposition 4.4 (Cf. van Hee et al. [38, Theorem 3.5], van der Wal [37, Theorem 3.3])**.** *The following holds:*

$$\limsup_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x), \quad \forall\, x \in S_0, \qquad \lim_{n\to\infty} T^n(\mathbf{0})(x) = J^*(x), \quad \forall\, x \in S_{0^+},$$

where the sets S_0 and S_{0^+} are given by

$$S_0 = \Big\{ x \in S \,\Big|\, \liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{J^*(x_n)\} \leq 0 \Big\}, \qquad S_{0^+} = \Big\{ x \in S \,\Big|\, \limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{J^*(x_n)\} \leq 0 \Big\}.$$

In particular, if $\limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\big\{J^*(x_n)\big\} \leq 0$ for all $x \in S$, then $T^n(\mathbf{0}) \to J^*$.

Proof. Consider any policy $\bar{\pi}$ and state x. Let $J_{\bar{\pi},n}(x)$ be the n-stage cost of $\bar{\pi}$ for the initial state x. Since $T^n(J^*) = J^*$, by the definition of T^n, we have

$$J^*(x) = T^n(J^*)(x) \leq J_{\bar{\pi},n}(x) + \mathbb{E}_x^{\bar{\pi}}\big\{J^*(x_n)\big\} \leq J_{\bar{\pi},n}(x) + \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\big\{J^*(x_n)\big\},$$

so, by minimizing the right-hand side over $\bar{\pi}$, we have

$$J^*(x) \leq T^n(\mathbf{0})(x) + \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\big\{J^*(x_n)\big\}.$$

By taking liminf over n in the right-hand side of the preceding inequality, it follows that

$$J^*(x) \leq \limsup_{n\to\infty}\, T^n(\mathbf{0})(x) + \liminf_{n\to\infty}\, \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\big\{J^*(x_n)\big\},$$
$$J^*(x) \leq \liminf_{n\to\infty}\, T^n(\mathbf{0})(x) + \limsup_{n\to\infty}\, \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\big\{J^*(x_n)\big\}.$$

Using the definitions of the sets S_0, S_{0^+}, we then have that $\limsup_{n\to\infty} T^n(\mathbf{0})(x) \geq J^*(x)$ for $x \in S_0$ and $\liminf_{n\to\infty} T^n(\mathbf{0})(x) \geq J^*(x)$ for $x \in S_{0^+}$. Since $\limsup_{n\to\infty} T^n(\mathbf{0})(x) \leq J^*(x)$ for all x, the proposition follows. □

Remark 4.2. For the (GC) model, similar to Lemma 3.2, it holds[18] that for any state $x \in S$ and policy π such that $J_\pi(x) < +\infty$,

$$\limsup_{n\to\infty}\, \mathbb{E}_x^{\pi}\big\{J^*(x_n)\big\} \leq 0. \tag{4.24}$$

We can compare Eq. (4.24) with the sufficient condition for $T^n(\mathbf{0})(x) \to J^*(x)$ given by the definition of S_{0^+} in Prop. 4.4: $\limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\{J^*(x_n)\} \leq 0$, which not only implies Eq. (4.24) but also requires that as n increases, the expectations $\mathbb{E}_x^{\pi}\big\{J^*(x_n)\big\}$ under different policies π must approach the interval $(-\infty, 0]$ at "comparable speeds." □

Remark 4.3. By Lemma 3.1(b), $\mathbb{E}_x^{\pi}\{\sum_{k=n}^{\infty} g(x_k,u_k)\} \geq \mathbb{E}_x^{\pi}\{J^*(x_n)\}$ for any policy π, state x and $n \geq 0$. So Prop. 4.4 implies the following "tail conditions" for convergence of value iteration, which can be relatively easier to verify in practice than the conditions of Prop. 4.4:

$$\limsup_{n\to\infty}\, T^n(\mathbf{0})(x) = J^*(x), \qquad \forall\, x \text{ s.t. } \liminf_{n\to\infty}\, \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\left\{\sum_{k=n}^{\infty} g(x_k,u_k)\right\} \leq 0, \tag{4.25}$$

$$\lim_{n\to\infty}\, T^n(\mathbf{0})(x) = J^*(x), \qquad \forall\, x \text{ s.t. } \limsup_{n\to\infty}\, \sup_{\pi\in\Pi} \mathbb{E}_x^{\pi}\left\{\sum_{k=n}^{\infty} g(x_k,u_k)\right\} \leq 0. \tag{4.26}$$

The statement in Eq. (4.26) is a "pointwise version" of a theorem in Hernández-Lerma and Lasserre [20, Theorem 9.3.5(b)].[19] □

[18] The proof of Eq. (4.24) is similar to that of Lemma 3.2: By Lemma 3.1(b), we have

$$J_\pi(x) = \mathrm{E}_x^{\pi}\left\{\sum_{k=0}^{n-1} g(x_k,u_k)\right\} + \mathrm{E}_x^{\pi}\left\{\sum_{k=n}^{\infty} g(x_k,u_k)\right\} \geq \mathrm{E}_x^{\pi}\left\{\sum_{k=0}^{n-1} g(x_k,u_k)\right\} + \mathrm{E}_x^{\pi}\left\{J^*(x_n)\right\}.$$

Since $+\infty > J_\pi(x) = \lim_{n\to\infty} \mathrm{E}_x^{\pi}\left\{\sum_{k=0}^{n-1} g(x_k,u_k)\right\} > -\infty$ by the (GC) condition (2.1), Eq. (4.24) follows.

[19] The statement of [20, Theorem 9.3.5(b)] differs from Eq. (4.26) here in that it has "liminf" instead of "limsup" in the condition on x. However, an inspection of its proof shows that the correct version should be as given in Eq. (4.26). Indeed, otherwise, Example 5.1 (Section 5) would serve as a counterexample: there the condition $\liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathrm{E}_x^{\pi}\left\{\sum_{k=n}^{\infty} g(x_k,u_k)\right\} \leq 0$ is met at all states x, and yet $\{T^n(\mathbf{0})\}$ does not have a limit.

The Special Case (UD)

We consider now the β-discounted problems in the (UD) class (cf. Def. 2.2). For these problems, the dynamic programming operators T, T_+, T_-, T_0 are defined as for the (GC) model, except that the integrals in their definitions are now multiplied by the discount factor β (cf. the discussion preceding Cor. 3.2). As mentioned before Def. 2.2, it is a standard procedure to reformulate these discounted problems as undiscounted total cost problems on an augmented state space $S \cup \{\Delta\}$ with Δ representing an absorbing, cost-free, termination state. By first applying Theorems 4.2, 4.3 to the equivalent undiscounted problems, and then expressing the results in terms of the original discounted problems, we obtain the following version of Theorems 4.2 and 4.3 for the unbounded discounted model (UD).

Corollary 4.8. (UD) *The conclusions of Theorems 4.2, 4.3 hold, under the same conditions of these theorems, if the sets S_{0^+}, S_0 in their statements are redefined as*

$$S_{0^+} := \Big\{x \in S \;\Big|\; \limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\beta^n \phi(x_n)\} \le 0\Big\},$$
$$S_0 := \Big\{x \in S \;\Big|\; \liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\beta^n \phi(x_n)\} \le 0\Big\}.$$

Corollary 4.9. (UD) *Suppose $\sup_{x\in S} \sup_{\pi\in\Pi} J_\pi^-(x) < \infty$. If $T_+^{\bar n}(\mathbf{0}) \ge \alpha J^* + b$ for some $\bar n \ge 0$, $\alpha \in (0,1]$, and scalar $b \le 0$ (which holds if $T_+^{\bar n}(\mathbf{0}) \ge \alpha J_\pi + b$ for some policy π), then $T^n(\mathbf{0}) \to J^*$.*

Proof. Using the assumption and the first inequality in (4.11), we have

$$T^{\bar n}(\mathbf{0}) \ge T_+^{\bar n}(\mathbf{0}) + T_-^{\bar n}(\mathbf{0}) \ge \alpha J^* + b + T_-^{\bar n}(\mathbf{0}) \ge \alpha J^* + \bar b,$$

where $\bar b$ in the last expression is given by $\bar b = b - \sup_{x\in S} \sup_{\pi\in\Pi} J_\pi^-(x) > -\infty$, and the last inequality holds because $T_-^{\bar n}(\mathbf{0})(x) \ge -\sup_{\pi\in\Pi} J_\pi^-(x)$ for all $x \in S$. Let $\phi(\cdot) \equiv \bar b$ on S. Since for all states x and policies π, $\mathbb{E}_x^\pi\{-\beta^n \phi(x_n)\} \le \beta^n |\bar b| \to 0$, the desired result follows from Cor. 4.8. □

Remark 4.4. The condition $T_+^{\bar n}(\mathbf{0}) \ge \alpha J^* + b$ in Cor. 4.9 can also be compared to the bridging condition, since it relates the infinite-horizon cost J^* to the finite-horizon cost $T_+^{\bar n}(\mathbf{0})$, albeit $T_+^{\bar n}(\mathbf{0})$ is from a different problem, the one associated with g_+. Note that this condition does *not* require the convergence of value iteration $T_+^n(\mathbf{0})$ to J^{*+}. □

5 Illustrative Examples

We start with two examples, which are modified versions of [15, Examples 6.6, 6.11]. They are counterexamples to the convergence of the value iteration $T^n(\mathbf{0})$ for the (GC) model.

Example 5.1 (The limit of $\{T^n(\mathbf{0})\}$ does not exist)**.** The problem is deterministic. The state space is a subset of nonnegative integer pairs: $S = \{(0,0)\} \cup \big(\{1,2,\ldots\} \times \{0,1,2,\ldots\}\big)$. For any state (i,j) with $i > 0$, only control 0 can be applied, under which the system moves to state $(i+1, j)$. At state $(0,0)$, the feasible control set is $\{0,1,2,\ldots\}$, and with a feasible control k, the system moves to state $(1,k)$. All the one-stage costs are zero, except that

$$g\big((2i+2,i),0\big) = 2, \qquad g\big((2i+1,i),0\big) = g\big((2i+3,i),0\big) = -1, \qquad i = 0,1,\ldots.$$

At state $(0,0)$, the optimal cost $J^*\big((0,0)\big) = 0$, but

$$\liminf_{n\to\infty} T^n(\mathbf{0})\big((0,0)\big) = -1, \qquad \limsup_{n\to\infty} T^n(\mathbf{0})\big((0,0)\big) = 0.$$

Thus, $\{T^n(\mathbf{0})\}$ does not have a limit, although $\limsup_{n\to\infty} T^n(\mathbf{0}) = J^*$, as can be verified. □

Example 5.2 ($J_\infty \neq J^*$)**.** The state space, feasible control sets, and the dynamics of the system are as given in Example 5.1. All the one-stage costs are zero, except that

$$g\big((i,i),0\big) = 1, \qquad i = 0,1,2,\ldots.$$

Then at state $(0,0)$, $T^n(\mathbf{0})\big((0,0)\big) = 0$ for all $n \geq 0$, but $J^*\big((0,0)\big) = 1$. □

The above examples have a countable state space. It is well-known, however, that even for finite state and control problems, $J_\infty \neq J^*$ can happen, as Example 4.1 showed.

We now use these examples to illustrate the convergence of value iteration from above stated in Theorem 3.1, and the partial convergence of value iteration stated in Theorems 4.2, 4.3 or their corollaries.

Example 5.3 (Example 4.1 Cont'd)**.** Consider Example 4.1. It has 3 states $\{0,1,2\}$. Let us write a cost function J in vector form as $J = (J(0), J(1), J(2))$. In this problem any policy π is optimal, with $J_\pi = J^* = (0,1,0)$, whereas $J_\infty = T(\mathbf{0}) = (0,1,-1)$ and it is a fixed point of T. Let us apply Theorem 4.2 (or equivalently, Cor. 4.6 or Theorem 4.3) with $\bar{n} = 1, \alpha = 1$ and $\phi = T(\mathbf{0}) - J_\pi = (0,0,-1)$. Starting from state $x = 2$, there is a policy to keep the system forever at state 2, so we have

$$\liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} = \limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} = 1, \qquad x = 2.$$

Starting from other states x, the above liminf and limsup expressions are equal to zero. So the sets S_{0^+}, S_0 in Theorem 4.2 are given by $S_{0^+} = S_0 = \{0,1\}$, and according to Theorem 4.2 or Cor. 4.6, for $x \in \{0,1\}$, $T^n(\mathbf{0})(x)$ must converge to $J^*(x)$, which is what we observe in this example.

Consider now value iteration starting from $J \geq J^*$. We have $J^{*+} = (0,1,0), J^{*-} = (0,0,0)$, so according to Theorem 3.1, $T^n(J) \to J^*$ for any $J = (0,c,0)$, $c \geq 1$. Indeed we have $T(J) = (0,1,0) = J^*$ in this example. □

Example 5.4 (Example 5.1 Cont'd)**.** Consider Example 5.1. In this problem any policy π is optimal. Let us apply Theorem 4.2 or Cor. 4.6 with $\bar{n} = 0, \alpha = 1$ and the function $\phi = -J_\pi = -J^*$. In particular, we have

$$J^*\big((k,i)\big) = \begin{cases} 1 & k = 2i+2; \\ -1 & k = 2i+3; \\ 0 & \text{otherwise;} \end{cases}$$

so

$$\phi\big((k,i)\big) = 0 \ \text{ if } k \geq 2i+4; \qquad \phi\big((k,i)\big) < 0 \quad \text{if and only if } k = 2i+2. \tag{5.1}$$

Starting from any state $x = (k,i) \neq (0,0)$, the states evolve as $(k+1,i), (k+2,i), \ldots$, so

$$\liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} = \limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} = 0, \qquad x \neq (0,0).$$

For state $x = (0,0)$, we have $\limsup_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} = 1$ because for any positive even number n, there exists a policy π with $\mathbb{E}_x^\pi\{-\phi(x_n)\} = \mathbb{E}_x^\pi\{J^*(x_n)\} = 1$. On the other hand, in view of the second relation in Eq. (5.1), under any policy π, we have $\mathbb{E}_x^\pi\{-\phi(x_n)\} \leq 0$ if n is an odd number, so

$$\liminf_{n\to\infty} \sup_{\pi\in\Pi} \mathbb{E}_x^\pi\{-\phi(x_n)\} = 0, \qquad x = (0,0).$$

Thus, the sets S_{0^+}, S_0 in Theorem 4.2 are given by

$$S_{0^+} = S \setminus \{(0,0)\}, \qquad S_0 = S.$$

Since $S_0 = S$, according to Theorem 4.2 or Cor. 4.6, $\limsup_{n\to\infty} T^n(\mathbf{0}) = J^*$, which is consistent with what actually happens in this example. According to Theorem 4.2 or Cor. 4.6, $T^n(\mathbf{0})(x) \to J^*(x)$ for all $x \in S_{0^+}$, i.e., for all $x \neq (0,0)$, which is also consistent with what we see in the example.

Consider now the convergence of value iteration from above. Let us calculate J^{*+}, J^{*-}:

$$J^{*+}\big((k,i)\big) = \begin{cases} 2 & k \le 2i+2; \\ 0 & \text{otherwise;} \end{cases} \qquad J^{*-}\big((k,i)\big) = \begin{cases} 2 & k \le 2i+1; \\ 1 & k = 2i+2 \text{ or } 2i+3; \\ 0 & \text{otherwise.} \end{cases}$$

Comparing $J^{*+} + J^{*-}$ with the optimal cost function J^* calculated earlier, we see that according to Theorem 3.1, if we start value iteration, for instance, with the function J given by

$$J\big((k,i)\big) = \begin{cases} c & k \le 2i+2; \\ -1 & k = 2i+3; \\ 0 & \text{otherwise;} \end{cases} \qquad \text{where } \ c \ge 1,$$

then we should have $T^n(J) \to J^*$. Indeed this is the case; in particular, for state $(0,0)$, we have that $T^n(J)\big((0,0)\big) = 0 = J^*\big((0,0)\big)$ for all $n \ge 3$.

As discussed immediately after Theorem 3.1, we can also start value iteration from $J^{*+} - J^{*-}$ to have monotonic convergence: $T^n(J^{*+} - J^{*-}) \downarrow J^*$. It can be seen that this is the case here; actually, $J^* = J^{*+} - J^{*-}$ in this example. □

6 Concluding Remarks

In this paper we considered a general class of total cost problems under the (GC) condition, and we presented several convergence results that characterize the behavior of value iteration for this class. Our result on convergence of value iteration from above holds for all problems in the (GC) class. It is also relevant to the policy iteration method, since policy iteration works with functions lying above the optimal cost function, and it sheds some light on why policy iteration can get stuck at a suboptimal policy in the total cost case (cf. [28, Example 7.3.4]). In the context of finding near optimal policies, another subject to which this result may be potentially relevant is finding $\epsilon\ell$-optimal stationary policies, where ℓ is a given nonnegative function on S (see e.g., [14, 18, 27, 31]; see also the book [28, Chap. 7.2.3], the survey [15, Chap. 6.7-6.9] and the references therein); investigation of this topic is needed. Regarding value iteration, a question that requires further research is how to find efficiently an initial function in the set with guaranteed convergence. We can start with the function J^{*+} or $J^{*+} - J^{*-}$, for example, but calculating J^{*+} or J^{*-} can be as difficult when the control space is infinite. We can also start with the cost function of some policy, but in general it can be hard to verify whether this function satisfies the convergence condition of our theorem. Another interesting question is whether by choosing suitably an initial function, one can not only ensure but also speed up the convergence of value iteration. (For average cost problems, Chen and Meyn [9] and Meyn [26] have shown that this is the case.)

We envision that our result on partial convergence of value iteration can be useful in practice in several ways. First, we may combine this result with the knowledge of a given problem to ensure that $T^n(\mathbf{0}) \to J^*$ holds in the given problem, or to find a subset of states for which the convergence holds, thereby extracting out useful information from the iterates produced by value iteration. Second, if we know a subset C of states x on which $\lim_{n\to\infty} T^n(\mathbf{0})(x)$ or $\limsup_{n\to\infty} T^n(\mathbf{0})(x)$ equals $J^*(x)$, we may estimate these limit values and then restart value iteration by setting the values of the initial function on C according to the estimates. This heuristic way of applying value iteration is in the spirit of transfinite value iteration, and may turn out to have some interesting properties upon further study.

In summary, the analytical results presented in this paper can provide useful guidance or background, we believe, for the application of the value iteration method in practice, as well as for further theoretical or computational study of this method. As part of the theoretical foundations, these results can also be useful for applying learning and simulation-based approaches [4, 26, 34, 35] and stochastic approximation techniques [8, 23] to solve total cost MDP. In that context, a particular subject for future research is Q-learning [36, 39] and related asynchronous stochastic value iteration algorithms, and their convergence properties for total cost MDP under the (GC) condition.

Acknowledgments

Several ideas presented in this paper grew out of a previous work [41] that I did with Professor Dimitri Bertsekas, at the Laboratory for Information and Decision Systems (LIDS), MIT. I thank Professor Bertsekas for helpful discussions then and during the course of writing this paper. I also thank Professors Eugene Feinberg and William Sudderth for helpful comments; Professor Sean Meyn for pointing out the material on value iteration in his book [26], which motivated the first convergence result given in this paper; and two anonymous reviewers, whose feedback helped to improve the presentation of this paper. This work was completed after I joined the Reinforcement Learning Artificial Intelligence (RLAI) research group at the Department of Computing Science, University of Alberta, and I am very thankful for the support from Alberta Innovates – Technology Futures and from the RLAI group for my research.

References

[1] D. P. Bertsekas. Monotone mappings with application in dynamic programming. *SIAM J. Control Opt.*, 15:438–464, 1977.

[2] D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*. Academic Press, New York, 1978.

[3] D. P. Bertsekas and J. N. Tsitsiklis. An analysis of stochastic shortest path problems. *Math. Oper. Res.*, 16:580–595, 1991.

[4] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, 1996.

[5] D. Blackwell. Positive dynamic programming. In *Proc. 5th Berkeley Sympos. Math. Satist. and Probability*, pages 415–418, 1965.

[6] D. Blackwell. A Borel set not containing a graph. *Ann. Math. Statist.*, 39:1345–1347, 1968.

[7] D. Blackwell, D. Freedman, and M. Orkin. The optimal reward operator in dynamic programming. *Ann. Probability*, 2:926–941, 1974.

[8] V. S. Borkar. *Stochastic Approximation: A Dynamic Viewpoint*. Hindustan Book Agency, New Delhi, 2008.

[9] R. R. Chen and S. Meyn. Value iteration and optimization of multiclass queueing networks. *Queueing Systems*, 32:65–97, 1999.

[10] R. M. Dudley. *Real Analysis and Probability*. Cambridge University Press, Cambridge, 2002.

[11] F. Dufour, M. Horiguchi, and A. B. Piunovskiy. The expected total cost criterion for Markov decision processes under constraints: a convex analytic approach. *Adv. Appl. Prob.*, 44:774–793, 2012.

[12] F. Dufour and A. B. Piunovskiy. The expected total cost criterion for Markov decision processes under constraints. *Adv. Appl. Prob.*, 45:837–859, 2013.

[13] E. B. Dynkin and A. A. Yushkevich. *Controlled Markov Processes*. Springer, New York, 1979.

[14] E. A. Feinberg. Sufficient classes of strategies in discrete dynamic programming II: locally stationary strategies. *SIAM Theory Prob. Appl.*, 32:478–493, 1987.

[15] E. A. Feinberg. Total reward criteria. In E. A. Feinberg and A. Shwartz, editors, *Handbook of Markov Decision Processes*. Springer, New York, 2002.

[16] E. A. Feinberg, P. O. Kasyanov, and N. V. Zadoianchuk. Average cost Markov decision processes with weakly continuous transition probabilities. *Math. Oper. Res.*, 37:591–607, 2012.

[17] E. A. Feinberg, P. O. Kasyanov, and M. Z. Zgurovsky. Convergence of value iterations for total-cost MDPs and POMDPs with general state and action sets. In *Proc. IEEE International Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, 2014.

[18] E. A. Feinberg and I. M. Sonin. Stationary and Markov policies in countable state dynamic programming. In *Lecture Notes in Math. 1021*, pages 111–129. Springer-Verlag, Berlin and New York, 1983.

[19] R. Hartley. A simple proof of Whittle's bridging condition in dynamic programming. *J. Appl. Prob.*, 17:1114–1116, 1980.

[20] O. Hernández-Lerma and J. B. Lasserre. *Further Topics on Discrete-Time Markov Control Processes*. Springer, New York, 1999.

[21] D. M. Kreps and E. L. Porteus. On the optimality of structured policies in countable stage decision processes. II: positive and negative problems. *SIAM J. Appl. Math.*, 32:457–466, 1977.

[22] K. Kuratowski. *Topology I*. Academic Press, New York, 1966.

[23] H. J. Kushner and G. G. Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, New York, 2nd edition, 2003.

[24] A. Maitra and W. Sudderth. The optimal reward operator in negative dynamic programming. *Math. Oper. Res.*, 17:921–931, 1992.

[25] A. Maitra and W. Sudderth. *Discrete Gambling and Stochastic Games*. Springer, New York, 1996.

[26] S. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, Cambridge, 2008.

[27] D. Ornstein. On the existence of stationary optimal strategies. *Proc. Am. Math. Soc.*, 20:563–569, 1969.

[28] M. L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming*. John Wiley & Sons, New York, 1994.

[29] M. Schäl. Conditions for optimality in dynamic programming and for the limit of n-stage optimal policies to be optimal. *Z. Wahrscheinlichkeitstheorie verw. Gebiete*, 32:179–196, 1975.

[30] M. Schäl. Stationary policies in dynamic programming models under compactness assumptions. *Math. Oper. Res.*, 8:366–372, 1983.

[31] M. Schäl and W. Sudderth. Stationary policies and Markov policies in Borel dynamic programming. *Probab. Th. Rel. Fields*, 74:91–111, 1987.

[32] S. E. Shreve. Resolution of measurability problems in discrete-time stochastic control. In *Stochastic Control Theory and Stochastic Differential Systems*, pages 580–587. Springer, Berlin, 1979.

[33] S. E. Shreve and D. P. Bertsekas. Universally measurable policies in dynamic programming. *Math. Oper. Res.*, 4:15–30, 1979.

[34] R. S. Sutton and A. G. Barto. *Reinforcement Learning*. MIT Press, Cambridge, 1998.

[35] C. Szepesvári. *Algorithms for Reinforcement Learning*. Morgan & Claypool, 2010.

[36] J. N. Tsitsiklis. Asynchronous stochastic approximation and Q-learning. *Mach. Learn.*, 16:185–202, 1994.

[37] J. van der Wal. *Stochastic Dynamic Programming*. The Mathematical Centre, Amsterdam, 1981.

[38] K. M. van Hee, A. Hordijk, and J. van der Wal. Successive approximations for convergent dynamic programming. In H. Tijms and J. Wessels, editors, *Markov Decision Theory*, pages 183–211. The Mathematical Centre, Amsterdam, 1977.

[39] C. J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge Univ., England, 1989.

[40] P. Whittle. A simple condition for regularity in negative programming. *J. Appl. Prob.*, 16:305–318, 1979.

[41] H. Yu and D. P. Bertsekas. A mixed value and policy iteration method for stochastic control with universally measurable policies. LIDS Technical Report 2905, MIT, 2013. Revised 2014; to appear in *Math. Oper. Res.* Online at: `http://arxiv.org/abs/1308.3814`.

[42] M. Zinsmeister. Les derivations analytiques. In *Seminaire de Probabilities XXIII*, Lecture Notes in Math. 1372, pages 21–46. Springer-Verlag, Berlin and New York, 1989.

Appendices

A Some Basic Properties of the (GC) Model in the Universal Measurability Framework

In this appendix, we prove some basic properties about policy costs for the (GC) total cost model, in the universal measurability framework. First, we prove Lemma 3.1, which we recall here for convenience.

Lemma 3.1 (Two basic properties)**.**

(a) For any given state x and policy π, there exists a Markov policy π' with

$$J^+_{\pi'}(x) = J^+_\pi(x), \qquad J^-_{\pi'}(x) = J^-_\pi(x), \qquad J_{\pi'}(x) = J_\pi(x).$$

(b) For any state x, policy π, and $n \geq 0$,

$$\mathbb{E}^\pi_x\left\{\sum_{k=n}^\infty g(x_k,u_k)\right\} \geq \mathbb{E}^\pi_x\left\{J^*(x_n)\right\}, \qquad \mathbb{E}^\pi_x\left\{\sum_{k=n}^\infty g_-(x_k,u_k)\right\} \geq \widetilde{\mathbb{E}}^\pi_{n,x}\left\{J^{*-}(x_n)\right\}.$$

Proof. (a) The proof arguments of [2, Prop. 8.1] are applicable here. They show that for any given policy π, one can construct a Markov policy π' such that for every $k \geq 0$, the marginal distribution of the state and control variables (x_k, u_k) under π' is the same as that under π for the given initial state x. It then follows that $J^+_{\pi'}(x) = J^+_\pi(x)$ and $J^-_{\pi'}(x) = J^-_\pi(x)$. Since the difference between these two terms defines the total cost of π' and π at state x, we have $J_{\pi'}(x) = J_\pi(x)$.

(b) The statement is clearly true for a Markov policy π by the Markov property and the optimality of J^*. Now suppose π is a history-dependent policy. Let π' be a Markov policy used in the proof of part (a), for the given state x and policy π. Then, since the marginal distributions of (x_k, u_k) are the same under π and π' for all $k \geq 0$, we have that

$$\mathbb{E}^\pi_x\left\{\sum_{k=n}^\infty g(x_k,u_k)\right\} = \mathbb{E}^{\pi'}_x\left\{\sum_{k=n}^\infty g(x_k,u_k)\right\} \qquad \text{and} \qquad \mathbb{E}^\pi_x\left\{J^*(x_n)\right\} = \mathbb{E}^{\pi'}_x\left\{J^*(x_n)\right\}.$$

Since $\mathbb{E}^{\pi'}_x\{\sum_{k=n}^\infty g(x_k,u_k)\} \geq \mathbb{E}^{\pi'}_x\{J^*(x_n)\}$, we obtain the first relation in part (b).

To show the second relation in part (b), consider again the preceding Markov policy π' and write it as $\pi' = (\mu'_0, \mu'_1, \ldots)$. Denote $\pi'_n = (\mu'_n, \mu'_{n+1}, \ldots)$ (which is also a Markov policy). Similarly to the preceding proof, because the marginal distributions of (x_k, u_k) are the same under π and π' for all $k \geq 0$, we have

$$\mathbb{E}^\pi_x\left\{\sum_{k=n}^\infty g_-(x_k,u_k)\right\} = \mathbb{E}^{\pi'}_x\left\{\sum_{k=n}^\infty g_-(x_k,u_k)\right\} = \mathbb{E}^{\pi'}_x\left\{J^-_{\pi'_n}(x_n)\right\}$$

and

$$\widetilde{\mathbb{E}}^{\pi}_{n,x}\left\{J^{*-}(x_n)\right\} = \widetilde{\mathbb{E}}^{\pi'}_{n,x}\left\{J^{*-}(x_n)\right\}.$$

By the optimality of J^{*-}, $J^{*-} \leq J^{-}_{\pi'_n} \in \mathcal{M}(S)$, and therefore, by the definition (3.1) of outer integral,

$$\mathbb{E}^{\pi'}_{x}\left\{J_{\pi'_n}(x_n)\right\} \geq \widetilde{\mathbb{E}}^{\pi'}_{n,x}\left\{J^{*-}(x_n)\right\}.$$

Combining the preceding three relations, we obtain $\mathbb{E}^{\pi}_{x}\left\{\sum_{k=n}^{\infty} g_-(x_k,u_k)\right\} \geq \widetilde{\mathbb{E}}^{\pi}_{n,x}\left\{J^{*-}(x_n)\right\}$. □

Let us also verify another basic fact about policy costs in the (GC) model:

Lemma A.1. *For any stationary policy μ, $J_\mu = T_\mu(J_\mu)$.*

Proof. We can write $J_\mu = J^+_\mu - J^-_\mu$ and $T_\mu(J_\mu) = T^+_\mu(J^+_\mu) - T^-_\mu(J^-_\mu)$, where J^+_μ, T^+_μ (resp. J^-_μ, T^-_μ) are the total cost of μ and the dynamic programming operator associated with μ, respectively, in the total cost problem with nonnegative one-stage cost function g_+ (resp. g_-). By [2, Prop. 9.9], $J^+_\mu = T^+_\mu(J^+_\mu)$ and $J^-_\mu = T^-_\mu(J^-_\mu)$, the conclusion then follows. □

B Optimality Properties of the (GC) Model in the Universal Measurability Framework

In this appendix, we establish that in the universal measurability framework, the (GC) total cost model has the optimality properties stated by Theorems 2.1, 2.2. We recall the two theorems here for convenience:

Theorem 2.1 (Properties of the optimal cost function)**.**

(a) The optimal cost function J^* is lower semi-analytic.

(b) The optimality equation $J^* = T(J^*)$ holds.

Theorem 2.2 (Properties of optimal policies)**.**

(a) For any $\epsilon > 0$, there exists an ϵ-optimal nonrandomized semi-Markov policy.

(b) If for each state an optimal policy exists, then there exists an optimal randomized semi-Markov policy.

(c) If $J^* \geq 0$, then there exists an ϵ-optimal nonrandomized Markov policy; if in addition, for each state an optimal policy exists, then there exists an optimal nonrandomized stationary policy. Furthermore, if $J^* \geq 0$, a stationary policy μ is optimal if and only if $J^* = T_\mu(J^*)$.

Most of these optimality results can be derived by following the lines of analyses given in the paper [33] and the book [2, Chap. 9.2-9.3], although these works addressed infinite-horizon undiscounted problems with purely positive or negative one-stage costs only. We will not repeat the entire analyses here. In what follows, we first explain in Section B.1 the key arguments needed in order to set up the stage for proofs. We then outline the proofs for Theorem 2.1(a), Theorem 2.2(a) and Theorem 2.2(b) (see Sections B.2.1, B.2.2, respectively); and we give detailed proofs for Theorem 2.1(b), the part of Theorem 2.2(a) relating to the construction of a nonrandomized near-optimal policy, and Theorem 2.2(c) (see Sections B.2.3, B.2.4, B.2.5, respectively).

Let us start the discussion with the lower semi-analyticity property of J^* stated in Theorem 2.1(a). If $T^n(\mathbf{0}) \to J^*$, we can conclude immediately by [2, Lemma 7.30(2)] that J^* must be lower semi-analytic. Unfortunately, as in the positive model (P), value iteration need not converge for the (GC) model, so a direct proof that does not rely on the convergence of value iteration is needed. It is achieved by setting up a corresponding deterministic total cost problem [2, Chap. 9.2], to be referred to as (DM), on the space of marginal distributions of the state and control variables. The problem (DM), roughly speaking, is constructed to have the following properties:

(i) it is equivalent to the original problem in terms of policy costs; and

(ii) it allows direct applications of measurable selection results, in particular a Jankov-von Neumann type of selection theorem, which states that partial minimization of a lower semi-analytic function results in a lower semi-analytic function, together with a universally measurable minimizer or ϵ-minimizer [2, Props. 7.47, 7.50].

To prove the optimality properties of the original problem, we first analyze (DM) as an intermediate step, and we then transcribe the results from (DM) to the original problem, using the correspondences between the two problems.

B.1 A Deterministic Model and its Relation to the Original Problem

Recall that $\mathcal{P}(S)$ (resp. $\mathcal{P}(S\times C)$) denotes the set of probability measures on the Borel σ-algebra $\mathcal{B}(S)$ (resp. $\mathcal{B}(S\times C)$) of the state space S (resp. the state-control space $S\times C$). The graph of the control constraint U of a given total cost (GC) problem is the analytic set $\Gamma = \{(x,u) \mid x \in S, u \in U(x)\}$.

Definition B.1 ([2, Defs. 9.4-9.6])**.** For a total cost (GC) problem given in Section 2.1, the *corresponding deterministic control model (DM)* is defined as follows.

(i) The state space, control space, and model parameters are given by:

- State space $\mathcal{P}(S)$ and control space $\mathcal{P}(S\times C)$.
- Control constraint $\bar{U}$, which maps each state $p \in \mathcal{P}(S)$ to a set $\bar{U}(p)$ of feasible controls at p, defined as

$$\bar{U}(p) := \big\{\gamma \in \mathcal{P}(S\times C) \;\big|\; \gamma(B\times C) = p(B),\ \gamma(\Gamma) = 1,\, \forall\, B \in \mathcal{B}(S)\big\}.$$

  I.e., the controls at state p are those probability measures on $S\times C$ that have p as the marginal on S and assign probability one to the graph Γ of the original control constraint U. (Like U, the graph of $\bar{U}$ is an analytic set in $\mathcal{P}(S)\times\mathcal{P}(S\times C)$; see the proof of [2, Lemma 9.1].)

- System function $\bar{f} : \mathcal{P}(S\times C) \to \mathcal{P}(S)$, which maps each control $\gamma \in \mathcal{P}(S\times C)$ to a state $\bar{f}(\gamma) \in \mathcal{P}(S)$, defined according to the state transition kernel q of the original problem as

$$\bar{f}(\gamma)(B) := \int_{S\times C} q(B \,|\, x,u)\,\gamma\big(d(x,u)\big), \qquad B \in \mathcal{B}(S).$$

  This function $\bar{f}$ specifies how the states evolve in (DM).

- One-stage cost function $\bar{g} : \mathcal{P}(S\times C) \to [-\infty,+\infty]$, defined by the one-stage cost function g of the original problem as

$$\bar{g}(\gamma) := \bar{g}_+(\gamma) - \bar{g}_-(\gamma)$$

  where

$$\bar{g}_+(\gamma) := \int_{S\times C} g_+(x,u)\,\gamma\big(d(x,u)\big), \qquad \bar{g}_-(\gamma) := \int_{S\times C} g_-(x,u)\,\gamma\big(d(x,u)\big),$$

  with the convention $\infty - \infty = \infty$. (Like g, the functions $\bar{g}_+$, $-\bar{g}_-$ and $\bar{g}$ are all lower semi-analytic [2, Prop. 7.48 and Lemma 7.30].)

(ii) A *policy for (DM)* is a sequence of mappings $\bar{\pi} = (\bar{\mu}_0, \bar{\mu}_1, \ldots)$ such that for each $k \geq 0$, $\bar{\mu}_k : \mathcal{P}(S) \to \mathcal{P}(S\times C)$ and $\bar{\mu}_k(p) \in \bar{U}(p)$ for every $p \in \mathcal{P}(S)$ (i.e., $\bar{\mu}_k$ maps each state p to a feasible control at p). The set of all policies for (DM) is denoted by $\bar{\Pi}$.

(iii) Given an initial state $p_0 \in \mathcal{P}(S)$, applying a policy $\bar{\pi}$ in (DM) generates recursively a sequence of state and control paris $(p_0, \gamma_0), (p_1, \gamma_1), \ldots$ as

$$\gamma_k = \bar{\mu}_k(p_k), \qquad p_{k+1} = \bar{f}(\gamma_k), \quad k = 0, 1, \ldots, \tag{B.1}$$

where p_k is the state at time k, generated according to the system function $\bar{f}$.

We now discuss the relations between the original problem and the deterministic model (DM) just defined (cf. [2, Def. 9.9 and Prop. 9.2]):

(a) Given $p_0 \in \mathcal{P}(S)$ and a policy π of the original problem, let $p_k \in \mathcal{P}(S)$, $\gamma_k \in \mathcal{P}(S \times C)$, $k \geq 0$, be the marginal distributions of the state x_k, the state-control pairs (x_k, u_k), respectively, at time k, under the policy π, with the initial state distribution being p_0. We can define a corresponding policy $\bar{\pi} = (\bar{\mu}_0, \bar{\mu}_1, \ldots)$ for (DM) such that these marginal distributions γ_k, p_{k+1} are exactly the controls and states that would be generated in (DM) according to Def. B.1(iii), if $\bar{\pi}$ is applied and the initial state is p_0.

(b) Conversely, given $p_0 \in \mathcal{P}(S)$ and a policy $\bar{\pi}$ of (DM), let γ_k, p_{k+1}, $k \geq 0$ be the controls and states generated according to Def. B.1(iii) under $\bar{\pi}$, for the initial state p_0. Then we can use the decomposition result [2, Prop. 7.27 and Cor. 7.27.2] to define from the sequence $\{(p_k, \gamma_k)\}$ a corresponding Markov policy π for the original problem, such that these p_k, γ_k would be exactly the marginal distributions of x_k, (x_k, u_k), respectively, in the original problem if π is applied and the initial state distribution is p_0.

Given these correspondences between the two problems, under the (GC) condition (2.1) on the original problem, we can define the optimal total cost function $\bar{J}^*$ for (DM) as follows.

Definition B.2 (Cf. [2, Def. 9.6])**.** In (DM), for a state $p_0 \in \mathcal{P}(S)$, the *cost of a policy* $\bar{\pi}$ at p_0 is defined as

$$\bar{J}_{\bar{\pi}}(p_0) = \bar{J}^+_{\bar{\pi}}(p_0) - \bar{J}^-_{\bar{\pi}}(p_0)$$

(with the convention $\infty - \infty = \infty$), where

$$\bar{J}^+_{\bar{\pi}}(p_0) := \sum_{k=0}^{\infty} \bar{g}_+(\gamma_k), \qquad \bar{J}^-_{\bar{\pi}}(p_0) := \sum_{k=0}^{\infty} \bar{g}_-(\gamma_k),$$

and the sequence $\{\gamma_k\}$ is generated by $\bar{\pi}$ according to Def. B.1(iii). The *optimal cost* at p_0 is

$$\bar{J}^*(p_0) = \inf_{\bar{\pi} \in \bar{\Pi}} \bar{J}_{\bar{\pi}}(p_0).$$

For $x \in S$, denote by δ_x the Dirac measure that assigns probability 1 to the point x. From the correspondence between the two models described in (b) earlier, we see that if the original problem satisfies the (GC) condition (2.1), then in its corresponding (DM), for all $x \in S$ and policies $\bar{\pi}$ of (DM), $\bar{J}^-_{\bar{\pi}}(\delta_x)$ is finite and $\bar{J}_{\bar{\pi}}(\delta_x) > -\infty$, and moreover, $\bar{J}^*(\delta_x) > -\infty$.

Furthermore, by the relations described in (a)-(b) earlier, the total costs of the policies in the two problems correspond as well:

(a′) Given $x \in S$ and a policy π of the original problem, there exists a policy $\bar{\pi}$ of (DM) such that $J_\pi(x) = \bar{J}_{\bar{\pi}}(\delta_x)$.

(b′) Conversely, given $x \in S$ and a policy $\bar{\pi}$ of (DM), there exists a Markov policy π of the original problem such that $J_\pi(x) = \bar{J}_{\bar{\pi}}(\delta_x)$.

Thus we must have $J^*(x) = \bar{J}^*(\delta_x)$ for all $x \in S$.

Now a crucial property of (DM) is that its optimal cost function $\bar{J}^*$ is lower semi-analytic. This can be shown by writing $\bar{J}^*$ as the result of partial minimization of a lower semi-analytic function as follows. Let Δ consist of all points $(p_0, \gamma_0, \gamma_1, \ldots)$, called *admissible sequences* [2, Def. 9.7], such that

$$p_0 \in \mathcal{P}(S), \qquad (\gamma_0, \gamma_1, \ldots) \in \Delta_{p_0},$$

where Δ_{p_0} is the set of all control sequences $(\gamma_0, \gamma_1, \ldots)$ that can be generated by some policy of (DM) for the initial state p_0. Define a function $G : \Delta \to [-\infty, +\infty]$ by

$$G\big(p_0, \gamma_0, \gamma_1, \ldots\big) := \sum_{k=0}^{\infty} \bar{g}_+(\gamma_k) - \sum_{k=0}^{\infty} \bar{g}_-(\gamma_k)$$

(with $\infty - \infty = \infty$). Then Δ is an analytic subset of $\mathcal{P}(S) \times \big(\mathcal{P}(S \times C)\big)^\infty$ [2, Lemma 9.1]; the function G is lower semi-analytic [2, Lemma 7.30(4)]; and by the definition of the optimal cost function $\bar{J}^*$, it can be written equivalently as the result of partial minimization of G:

$$\bar{J}^*(p_0) = \inf_{(\gamma_0, \gamma_1, \ldots) \in \Delta_{p_0}} G\big(p_0, \gamma_0, \gamma_1, \ldots\big). \tag{B.2}$$

Hence, by [2, Prop. 7.47], $\bar{J}^*$ is lower semi-analytic.

Moreover, by a measurable selection theorem [2, Prop. 7.50], for any $\epsilon > 0$, we can select a measurable ϵ-minimizer for the optimization problem in (B.2). More precisely, there exists a universally measurable mapping $\psi : \mathcal{P}(S) \to \big(\mathcal{P}(S \times C)\big)^\infty$ such that for all $p_0 \in \mathcal{P}(S)$,

$$\psi(p_0) \in \Delta_{p_0} \quad \text{and} \quad G\big(p_0, \psi(p_0)\big) \le \begin{cases} \bar{J}^*(p_0) + \epsilon, & \text{if } \bar{J}^*(p_0) > -\infty; \\ -1/\epsilon, & \text{if } \bar{J}^*(p_0) = -\infty. \end{cases}$$

As noted earlier, $\bar{J}^*(\delta_x) > -\infty$ for all $x \in S$ since the original problem is a (GC) problem. Therefore,

$$\psi(\delta_x) \in \Delta_{\delta_x} \quad \text{and} \quad G\big(\delta_x, \psi(\delta_x)\big) \le \bar{J}^*(\delta_x) + \epsilon, \qquad \forall\, x \in S. \tag{B.3}$$

Furthermore, by [2, Prop. 7.50(b)], ψ can be chosen to attain the infimum in (B.2) whenever this is possible; i.e.,

$$G\big(p_0, \psi(p_0)\big) = \bar{J}^*(p_0), \qquad \forall\, p_0 \in D, \tag{B.4}$$

where

$$D = \Big\{ p \in \mathcal{P}(S) \;\Big|\; \exists\, (\gamma_0, \gamma_1, \ldots) \in \Delta_p \text{ with } \bar{J}^*(p) = G\big(p, \gamma_0, \gamma_1, \ldots\big) \Big\}. \tag{B.5}$$

B.2 Proof Arguments for Theorems 2.1, 2.2

B.2.1 J^* is Lower Semi-analytic

The correspondence between the original problem and (DM), as described in (a)-(b) and (a′)-(b′) earlier, implies that $J^*(x) = \bar{J}^*(\delta_x)$ for all $x \in S$. In other words, J^* is the composition of the function $\bar{J}^*$, which was just proved to be lower semi-analytic, with the function $x \mapsto \delta_x$, which is a homeomorphism [2, Cor. 7.21.1]. Then, since the composition $f_1 \circ f_2$ is lower semi-analytic for a lower semi-analytic function f_1 and a Borel-measurable function f_2 [2, Lemma 7.30(3)], we can conclude that J^* is a lower semi-analytic function, as stated in Theorem 2.1(a).

B.2.2 Existence of an ϵ-Optimal Semi-Markov Policy

We consider now the statements in Theorem 2.2 about the existence of ϵ-optimal or optimal policies. The measurable ϵ-minimizer ψ obtained in Eq. (B.3) for (DM) gives us the universally measurable mapping $\phi(x) = \psi(\delta_x)$ [2, Prop. 7.44], which maps $x \in S$ to a sequence of probability measures in $\big(\mathcal{P}(S \times C)\big)^\infty$. Using the decomposition result [2, Prop. 7.27] and using also the relations (a)-(b) and (a′)-(b′) between the two models mentioned earlier, one can construct from ϕ an ϵ-optimal, possibly randomized, semi-Markov policy for the original problem. The method of this construction is the same as the one used in the second half of the proof of [2, Prop. 9.20, p. 240]. That same part of the proof, together with Eqs. (B.3)-(B.5), also establishes Theorem 2.2(b) by constructing an optimal, possibly randomized, semi-Markov policy for the original problem.

We still have not shown that an ϵ-optimal policy can be chosen to be nonrandomized, as stated in Theorem 2.2(a). We shall prove this later in Appendix B.2.4.

B.2.3 The Optimality Equation $J^* = T(J^*)$

To show that the optimality equation $J^* = T(J^*)$ holds, we can establish first the optimality equation for the corresponding deterministic model (DM) and then transcribe the result from (DM) to the original problem. Alternatively, we can prove $J^* = T(J^*)$ directly, using the fact that there exists an ϵ-optimal policy for the original problem, which we just established. We give such a proof below. The line of analysis is similar to the one for [15, Theorem 6.3]. The existence of an ϵ-optimal policy is used to prove $J^* \leq T(J^*)$.

Proof of Theorem 2.1(b). We prove $J^* \leq T(J^*)$ first. Let $\epsilon > 0$. As we showed earlier, there exists a policy $\hat{\pi}$ which is ϵ-optimal for all states:

$$J_{\hat{\pi}}(x) \leq J^*(x) + \epsilon, \qquad \forall\, x \in S.$$

By the selection theorem [2, Prop. 7.50], there exists a nonrandomized stationary policy μ such that

$$T_\mu(J^*) \leq T(J^*) + \epsilon.$$

Then, for the policy $\pi = (\mu, \hat{\pi})$ (the concatenation of μ with $\hat{\pi}$), we have

$$J^* \leq J_\pi = T_\mu(J_{\hat{\pi}}) \leq T_\mu(J^*) + \epsilon \leq T(J^*) + 2\epsilon.$$

Since ϵ is arbitrary, we obtain $J^* \leq T(J^*)$.

We now prove $J^* \geq T(J^*)$. Consider an arbitrary state x. Take a policy $\hat{\pi} = (\hat{\pi}_0, \hat{\pi}_1, \ldots)$ with

$$J_{\hat{\pi}}(x) \leq J^*(x) + \epsilon.$$

We also have $J_{\hat{\pi}}(x) \geq T(J^*)(x)$, since

$$\begin{aligned}
J_{\hat{\pi}}(x) &= \mathbb{E}_x^{\hat{\pi}}\{g(x_0, u_0)\} + \mathbb{E}_x^{\hat{\pi}}\left\{\sum_{k=1}^{\infty} g(x_k, u_k)\right\} \\
&\geq \mathbb{E}_x^{\hat{\pi}}\{g(x_0, u_0)\} + \mathbb{E}_x^{\hat{\pi}}\{J^*(x_1)\} \\
&= \int_C \left\{ g(x,u) + \int_S J^*(x')\, q(dx' \mid x, u) \right\} \hat{\pi}_0(du \mid x) \\
&\geq \inf_{u \in U(x)} \left\{ g(x,u) + \int_S J^*(x')\, q(dx' \mid x, u) \right\} = T(J^*)(x),
\end{aligned}$$

where the first inequality follows from Lemma 3.1(b). Combining the two relations, we obtain $J^*(x) \geq T(J^*)(x) - \epsilon$. Since x and ϵ are arbitrary, we obtain $J^* \geq T(J^*)$. So by the first part of the proof, $J^* = T(J^*)$. □

B.2.4 Existence of an ϵ-Optimal Nonrandomized Semi-Markov Policy

We showed earlier that an ϵ-optimal, possibly randomized, semi-Markov policy exists for the (GC) model. We will now prove that this policy can be chosen to be nonrandomized, thus establishing Theorem 2.2(a) completely. The construction of this policy is similar to the one given in the proof of [2, Prop. 9.20, p. 239] for the negative costs model (N), except for a few details due to the difference between the negative costs model (N) and the (GC) model. To account for this difference, we will need in the proof the convergence property of value iteration in the (GC) model; in particular, we will use the fact that value iteration converges to J^* from above if it starts from J^{*+}, for instance. We give a complete proof below.

Proof of Theorem 2.2(a). Let $\epsilon > 0$ be given. Recall that J^{*+} is the optimal cost function of the positive costs problem with one-stage cost function g_+. For $k \geq 0$, let $J_k = T^k(J^{*+})$ and let

$$A_k = \big\{x \in S \mid J_k(x) \leq J^*(x) + \epsilon/4\big\}.$$

Since J_k, J^* are lower semi-analytic, the sets A_k are universally measurable, and since $J_k \downarrow J^*$ (Theorem 3.1), we have $A_0 \subset A_1 \subset \cdots$ and $S = \bigcup_{k\geq 0} A_k$. For each $k \geq 0$, consider the $(k+1)$-stage problem corresponding to $T^k(J^{*+})$, which has J^{*+} as the terminal cost function at the last stage. Applying the same argument as in the proof of [2, Prop. 8.3(F$^+$)] for finite-horizon problems, we have that there exists a nonrandomized Markov policy π whose k-stage cost function $J_{\pi,k}$ satisfies that

$$J_{\pi,k}(x) + \mathbb{E}_x^\pi\{J^{*+}(x_k)\} \leq J_k(x) + \epsilon/4, \qquad \forall\, x \in S.$$

For states in the set A_k, this implies the bound,

$$J_{\pi,k}(x) + \mathbb{E}_x^\pi\{J^{*+}(x_k)\} \leq J^*(x) + \epsilon/2, \qquad \forall\, x \in A_k. \tag{B.6}$$

Let us write this policy π, which is associated with k, as $\pi^{(k)} = (\pi_0^{(k)}, \pi_1^{(k)}, \ldots)$.

Now let $\pi^+ = (\pi_0^+, \pi_1^+, \ldots)$ be a nonrandomized $\epsilon/2$-optimal Markov policy for the positive costs problem with one-stage cost function g_+. (The existence of such a policy follows from [2, Prop. 9.19(P)].) Since $g \leq g_+$, the cost of π^+ in the original problem can be bounded by

$$J_{\pi^+}(x) \leq J^+_{\pi^+}(x) \leq J^{*+}(x) + \epsilon/2, \qquad \forall\, x \in S.$$

Consequently, on each set A_k, $k \geq 0$, we have, by Eq. (B.6), that for the associated policy $\pi^{(k)}$,

$$J_{\pi^{(k)},k}(x) + \mathbb{E}_x^{\pi^{(k)}}\big\{J_{\pi^+}(x_k)\big\} \leq J^*(x) + \epsilon, \qquad \forall\, x \in A_k. \tag{B.7}$$

Define a semi-Markov policy

$$\pi^\epsilon = \big(\mu_0(du_0 \mid x_0),\, \mu_1(du_1 \mid x_0, x_1),\, \ldots,\, \mu_n(du_n \mid x_0, x_n),\, \ldots\big)$$

as follows. Partition the state space S into a countable number of disjoint sets, $A_k \setminus A_{k-1}$, $k \geq 0$, with $A_{-1} = \emptyset$. For initial states in the set $A_k \setminus A_{k-1}$, let π^ϵ be identical to the concatenation of the first k-stage of the policy $\pi^{(k)}$ and the policy π^+; more precisely, for $x_0 \in A_k \setminus A_{k-1}$, let

$$\mu_0(du_0 \mid x_0) = \pi_0^{(k)}(du_0 \mid x_0), \quad \ldots, \quad \mu_{k-1}(du_{k-1} \mid x_0, x_{k-1}) = \pi_{k-1}^{(k)}(du_{k-1} \mid x_{k-1}),$$

and let

$$\mu_n(du_n \mid x_0, x_n) = \pi_n^+(du_n \mid x_n), \quad n \geq k.$$

Then π^ϵ is a universally measurable semi-Markov policy, and it is nonrandomized since π^+ and $\pi^{(k)}, k \geq 0$ are all nonrandomized. Furthermore, for $x \in A_k \setminus A_{k-1}$, by the Markov property,

$$J_{\pi^\epsilon}(x) = J_{\pi^{(k)},k}(x) + \mathbb{E}_x^{\pi^{(k)}}\big\{J_{\pi^+}(x_k)\big\} \leq J^*(x) + \epsilon,$$

where the inequality follows from Eq. (B.7). Therefore, $J_{\pi^\epsilon}(x) \leq J^*(x) + \epsilon$ for all $x \in S$; i.e., π^ϵ is ϵ-optimal. □

B.2.5 Existence of ϵ-Optimal Markov or Optimal Stationary Policies when $J^* \geq 0$

Finally, we verify the statements in Theorem 2.2(c) about the existence of ϵ-optimal Markov policies or optimal stationary policies when $J^* \geq 0$. The conclusions of Theorem 2.2(c) are indeed the same as those known to hold for the positive costs model, and the proofs for the two models are also almost the same (cf. the proof of [2, Prop. 9.19(P)]). For clarity, however, we include the proof below. It uses the optimality equation $J^* = T(J^*)$ and the assumption $J^* \geq 0$.

Proof of Theorem 2.2(c). Since $J^* \geq 0$ by assumption, $T(J^*)(x) = J^*(x) \geq 0 > -\infty$ for each $x \in S$. So by the selection theorem [2, Prop. 7.50], given any $\epsilon > 0$, there exist stationary nonrandomized policies $\mu_k, k \geq 0$, such that

$$T_{\mu_k}(J^*) \leq T(J^*) + 2^{-k-1}\epsilon = J^* + 2^{-k-1}\epsilon. \tag{B.8}$$

Consider the nonrandomized Markov policy $\pi = (\mu_0, \mu_1, \ldots)$. By the monotonicity of the mappings $T_{\mu_0}, T_{\mu_1}, \ldots$, we have that for $n \geq 1$,

$$\begin{aligned}
\big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_n}\big)(\mathbf{0}) &\leq \big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_n}\big)(J^*) \\
&\leq \big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_{n-1}}\big)(J^*) + 2^{-n-1}\epsilon \\
&\leq \big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_{n-2}}\big)(J^*) + (2^{-n} + 2^{-n-1})\epsilon \\
&\quad \vdots \\
&\leq J^* + \sum_{k=0}^{n} 2^{-k-1}\epsilon \leq J^* + \epsilon,
\end{aligned}$$

where the assumption $J^* \geq 0$ is used in the first inequality, and Eq. (B.8) is used in the subsequent inequalities. Since $J_\pi = \lim_{n\to\infty} \big(T_{\mu_0} \circ T_{\mu_1} \circ \cdots T_{\mu_n}\big)(\mathbf{0})$ in the (GC) model, it follows that π is an ϵ-optimal policy. Hence there exists an ϵ-optimal nonrandomized Markov policy if $J^* \geq 0$.

We now prove the second statement in Theorem 2.2(c). If there exists an optimal policy for each state, by Lemma 3.1(a), there exists for each given state $x \in S$, a Markov policy $\pi = (\mu_0, \mu_1, \ldots)$ that is optimal for x. Denoting $\pi' = (\mu_1, \mu_2, \ldots)$, we have by the Markov property and the monotonicity of T_{μ_0},

$$T_{\mu_0}(J^*)(x) \leq T_{\mu_0}(J_{\pi'})(x) = J_\pi(x) = J^*(x) = T(J^*)(x).$$

This shows that for each x, the infimum defining $T(J^*)(x)$ is attained. Hence, by the selection theorem [2, Prop. 7.50], there exists a (universally measurable) nonrandomized stationary policy μ satisfying $T_\mu(J^*) = J^*$. Using the assumption $J^* \geq 0$ and the monotonicity of T_μ, we obtain

$$J_\mu = \lim_{n\to\infty} T_\mu^n(\mathbf{0}) \leq \lim_{n\to\infty} T_\mu^n(J^*) = J^*, \tag{B.9}$$

so we must have $J_\mu = J^*$, i.e., μ is optimal. Since μ is nonrandomized and stationary, this establishes the desired result.

We now prove the last statement in Theorem 2.2(c). Equation (B.9) already shows that when $J^* \geq 0$, any stationary policy μ satisfying $T_\mu(J^*) = J^*$ is optimal. Conversely, if μ is an optimal stationary policy, then $J_\mu = J^*$, so by Lemma A.1, $T_\mu(J^*) = J^*$. This completes the proof. □

C About Transfinite Value Iteration and an Important Lemma for Proving Theorem 4.1

In this appendix we first give an informal introduction of ordinals, in connection with transfinite value iteration. We then focus on an important lemma needed to establish Theorem 4.1 on the convergence of transfinite value iteration in the case where the optimal cost function is nonnegative.

C.1 Ordinals and Transfinite Value Iteration

Ordinals are well-ordered sets chosen in such a way that each well-ordered set is order-isomorphic to one and only one ordinal; in other words, each ordinal represents an order type. We can linearly order these order types so that a more complicated ordering is represented by a larger ordinal. The book [10, Chaps. A.3] gives a good introduction of the theory of ordinals, and the informal explanations we give here are partly based on this reference.

Intuitively, we can construct ordinals starting from the smallest set, the empty set, as follows. Ordinals with finite elements, named after the number of their elements, are:

$$\text{ordinal } 0 : \emptyset, \quad \text{ordinal } 1 : \{\emptyset\}, \quad \text{ordinal } 2 : \big\{\emptyset, \{\emptyset\}\big\}, \quad \text{ordinal } 3 : \big\{\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}\big\}, \quad \ldots.$$

In general, each element of an ordinal is a set itself, the *successor ordinal* $\xi + 1$ of an ordinal ξ is given by the set $\xi \cup \{\xi\}$, and any ordinal is well-ordered by the relation $\in$. The preceding finite ordinals are order-isomorphic to the simple finite sets $\emptyset, \{1\}, \{1, 2\}, \ldots$, respectively, with the usual ordering. The union of any set of ordinals is also an ordinal. The first infinite ordinal, denoted by ω, is the union of the finite ordinals $1, 2, \ldots$, and order-isomorphic to the set of positive integers. The ordinal ω is equal to the union of its elements—such an ordinal is called a *limit ordinal*. It can be shown that an ordinal is either a limit ordinal or a successor of some ordinal ξ, which is its largest element. In other words, with the "+1" and union operations, increasingly larger ordinals are constructed successively. (This process goes on indefinitely; there is no largest ordinal.)

The transfinite value iteration discussed in Section 4.1 is a recursion defined on the well-ordered set $\{0, 1, \ldots, \omega_1\}$, with ω_1 being the smallest ordinal that has an uncountably infinite number of elements. Iterations associated with those limit ordinals involve taking various limits of the "past" value iteration iterates, roughly speaking, whereas iterations associated with those ordinals of the form $\xi + 1$ are similar to ordinary value iteration with the dynamic programming operator. For example, consider the transfinite value iteration (4.5)-(4.6) given in Section 4.1, which involves the mapping $\widetilde{T}(J) = \max\{J, T(J)\}$. Let J_0 be the initial function, and denote $J_\xi = \widetilde{T}^\xi(J_0)$ for an ordinal $\xi < \omega_1$. For $\xi < \omega$ (the first infinite ordinal), the iterates J_ξ coincide with the iterates produced by ordinary fixed point iteration with $\widetilde{T}$:

$$J_1 = \widetilde{T}(J_0), \quad \ldots, \quad J_{n+1} = \widetilde{T}(J_n), \ \ldots,$$

which form an increasing sequence of functions by the increasing property of $\widetilde{T}$, $\widetilde{T}(J) \geq J$. For the limit ordinal ω, J_ω can be expressed as

$$J_\omega = \widetilde{T}\Big(\lim_{n\to\infty} J_n\Big) = \widetilde{T}\Big(\sup_{n\geq 0} J_n\Big),$$

and for the ordinal $\omega + 1$, by the increasing property of $\widetilde{T}$, $J_{\omega+1}$ is simply given by

$$J_{\omega+1} = \widetilde{T}\big(J_\omega\big).$$

Ordinarily, we stop value iteration at the ordinal ω. But when the limit function $\lim_{n\to\infty} J_n$ does not equal J^* (i.e., ordinary value iteration fails to converge to J^*), then transfinite value iteration can be used to analyze the behavior of value iteration, and it can shed light on the nature of value iteration, as it does for the positive costs model [24].

C.2 A Lemma for Convergence of Transfinite Value Iteration

We now state and outline the proof of an important lemma needed to establish Theorem 4.1 on the convergence of transfinite value iteration. Recall that Theorem 4.1 states that $\widetilde{T}^{\omega_1}(\mathbf{0}) = J^*$ if $J^* \geq 0$, where $\widetilde{T}(J) = \max\{J, T(J)\}$. Let $J_{\omega_1} = \widetilde{T}^{\omega_1}(\mathbf{0})$. In proving Theorem 4.1, we used the following lemma:

Lemma C.1. *The function J_{ω_1} is lower semi-analytic and satisfies*

$$J_{\omega_1} = \widetilde{T}(J_{\omega_1}) \geq T(J_{\omega_1}).$$

Lemma C.1 can be established by using the proof arguments given in Maitra and Sudderth's analysis of transfinite value iteration for the positive costs model [24]. In the rest of this appendix, we will explain the main proof steps—we do not include a detailed proof because most arguments are the same as in [24]. (The paper [24] considers the reward maximization framework and the negative rewards model, which is equivalent to the positive costs model in the cost minimization framework we use. However, because of this difference, some of our notation and definitions below will appear to be different from those in [24].)

First, note that if for some ordinal $\xi < \omega_1$, $\widetilde{T}^{\xi}(\mathbf{0})$ is a fixed point of $\widetilde{T}$ so that $J_{\omega_1} = \widetilde{T}^{\xi}(\mathbf{0})$, then it follows immediately that J_{ω_1} is lower semi-analytic, since $\sup_n f_n$ for a sequence of lower semi-analytic functions f_n is a lower semi-analytic function [2, Lemma 7.30]. However, in general, J_{ω_1} is the pointwise supremum of an uncountable set of lower semi-analytic functions, so the preceding lemma does not follow easily from the argument we just mentioned. The line of analysis to establish the lemma is as follows.

We are going to apply a deep theorem from the descriptive set theory. This theorem concerns fixed points of a mapping Ψ from the power-set of a Borel space Y to the power-set of Y. More specifically, a fixed point of Ψ is a set $E \subset Y$ satisfying $E = \Psi(E)$. For a given set $A \subset Y$, consider the set A_{ω_1} defined by the transfinite recursion,

$$A_0 = A; \qquad A_{\xi} = \Psi\left(\bigcap_{\eta<\xi} A_{\eta}\right), \quad 0 < \xi < \omega_1; \qquad A_{\omega_1} = \bigcap_{\xi<\omega_1} A_{\xi}. \tag{C.1}$$

The theorem states that if the mapping Ψ is such that:

(i) Ψ is monotone (i.e., $\Psi(E_1) \subset \Psi(E_2)$ if $E_1 \subset E_2 \subset Y$),

(ii) $\Psi(E) \subset E$ for $E \subset Y$, and Ψ is *uniformly analytic* in the sense that for any Polish space Ω and analytic set $C \subset \Omega \times Y$, the set $\{(\omega, y) \in \Omega \times Y \mid y \in \Psi(C_\omega)\}$ is analytic (here $C_\omega = \{y \in Y \mid (\omega, y) \in C\}$),

then for any analytic set $A \subset Y$, the set A_{ω_1} obtained by the recursion (C.1) is analytic and satisfies $A_{\omega_1} = \Psi(A_{\omega_1})$, and is indeed the largest subset of A invariant under Ψ.

For a proof of the preceding theorem, see Zinsmeister [42], which also showed that the uniform analyticity condition in (ii) implies the monotonicity of Ψ for analytic sets. According to [24, 42], this theorem is a special case of a very general result of Moschovakis; see [24, 42] for related references.

To apply this theorem to our problem, we take $Y = S \times [0, +\infty]$, and we equate nonnegative functions with their epigraphs and equate the mapping $\widetilde{T}$ to a mapping that maps epigraphs to epigraphs. Specifically, for any function $J \geq 0$, we consider equivalently its epigraph $\mathrm{epi}(J)$ given by

$$\mathrm{epi}(J) = \big\{(x, z) \in S \times [0, +\infty] \;\big|\; J(x) \leq z\,\big\},$$

and we define a mapping Ψ on the power-set of $S \times [0, +\infty]$ in such a way that for a lower semi-analytic function $J \geq 0$,

$$\mathrm{epi}\,\big(\widetilde{T}(J)\big) = \Psi\big(\,\mathrm{epi}(J)\big). \tag{C.2}$$

In other words, Ψ is the mapping equivalent to $\widetilde{T}$, on the space of epigraphs of functions. However, to apply the theorem, we also need to define Ψ for all subsets in $S \times [0, +\infty]$, including the epigraphs of those non-measurable functions. This is achieved in [24] by using an outer-integral formulation to define Ψ. Specifically, for any set $E \subset S \times [0, +\infty]$, we define

$$\Psi(E) = E \bigcap \left\{(x, z) \in S \times [0, +\infty] \;\middle|\; \inf_{u \in U(x)} \left\{ g(x, u) + \big(q(\cdot \mid x, u) \times \lambda\big)^*\big(E^c\big) \right\} \leq z \right\}. \tag{C.3}$$

Here λ is the Lebesgue measure on $[0,+\infty]$; for $p \in \mathcal{P}(S)$, $(p \times \lambda)^*$ denotes the outer measure with respect to the product measure $p \times \lambda$ on $S \times [0,+\infty]$; and E^c denotes the complement of a set E: $E^c = (S \times [0,+\infty]) \setminus E$.

By its definition (C.3), Ψ is monotone and $\Psi(E) \subset E$. By [24, Lemma 3.1], Ψ is uniformly analytic.[20] The mapping Ψ also satisfies the desired equality (C.2): for a lower semi-analytic function $J \geq 0$, we have by Fubini's theorem,

$$\big(q(\cdot \mid x,u) \times \lambda\big)^*\big(\operatorname{epi}(J)^c\big) = \int_S J(x')\, q(dx' \mid x,u),$$

so

$$\inf_{u \in U(x)} \Big\{ g(x,u) + \big(q(\cdot \mid x,u) \times \lambda\big)^*\big(\operatorname{epi}(J)^c\big) \Big\} = T(J)(x)$$

and $(x,z) \in \Psi\big(\operatorname{epi}(J)\big)$ if and only if

$$J(x) \leq z \;\text{ and }\; T(J)(x) \leq z \quad\Longleftrightarrow\quad \max\big\{J(x),\, T(J)(x)\big\} \leq z \quad\Longleftrightarrow\quad (x,z) \in \operatorname{epi}\big(\widetilde{T}(J)\big).$$

Now we can apply the theorem from the descriptive set theory mentioned earlier. Let A_{ω_1} and A_ξ, $\xi < \omega_1$, be given by the transfinite recursion (C.1) with $A_0 = \operatorname{epi}(\mathbf{0}) = S \times [0,+\infty]$. Denote $J_\xi = \widetilde{T}^\xi(\mathbf{0})$ for $\xi \leq \omega_1$; in particular $J_0 = \mathbf{0}$. Similar to [24, Lemma 3.3], we show by transfinite induction that $A_{\omega_1} = \operatorname{epi}(J_{\omega_1})$. We have $A_0 = \operatorname{epi}(J_0)$. For any given $\xi < \omega_1$, suppose $A_\eta = \operatorname{epi}\big(J_\eta\big)$ for all $\eta < \xi$, and we have

$$\begin{aligned} A_\xi &= \Psi\left(\bigcap_{\eta<\xi} A_\eta\right) = \Psi\left(\bigcap_{\eta<\xi} \operatorname{epi}\big(J_\eta\big)\right) \\ &= \Psi\left(\operatorname{epi}\Big(\sup_{\eta<\xi} J_\eta\Big)\right) = \operatorname{epi}\left(\widetilde{T}\Big(\sup_{\eta<\xi} J_\eta\Big)\right) = \operatorname{epi}\big(J_\xi\big), \end{aligned}$$

where we used Eq. (C.2) in the second to last equality, and we used the definition of $\widetilde{T}^\xi(\mathbf{0})$ in the last equality [cf. Eq. (4.5)]. By transfinite induction [10, Theorem 1.3.1], this implies $A_\xi = \operatorname{epi}(J_\xi)$ for all $\xi < \omega_1$, and consequently,

$$A_{\omega_1} = \bigcap_{\xi<\omega_1} A_\xi = \bigcap_{\xi<\omega_1} \operatorname{epi}(J_\xi) = \operatorname{epi}\Big(\sup_{\xi<\omega_1} J_\xi\Big) = \operatorname{epi}\big(J_{\omega_1}\big),$$

where in the last equality we used the definition $\widetilde{T}^{\omega_1}(\mathbf{0}) = \sup_{\xi<\omega_1} J_\xi$ [cf. Eq. (4.6)].

Now by the theorem mentioned earlier, A_{ω_1} is analytic and $A_{\omega_1} = \Psi(A_{\omega_1})$. This means, first, that $\operatorname{epi}(J_{\omega_1}) = A_{\omega_1}$ is an analytic set, which is equivalent to that J_{ω_1} is lower semi-analytic (cf. [2, p. 186]). Second,

$$\operatorname{epi}\big(J_{\omega_1}\big) = \Psi\big(\operatorname{epi}(J_{\omega_1})\big) = \operatorname{epi}\big(\widetilde{T}(J_{\omega_1})\big),$$

where the second equality follows from Eq. (C.2). Since two functions are identical if and only if they have the same epigraph, from the preceding relation we obtain

$$J_{\omega_1} = \widetilde{T}(J_{\omega_1}) = \max\big\{J_{\omega_1}, T(J_{\omega_1})\big\} \geq T(J_{\omega_1}).$$

This establishes Lemma C.1.

[20] More specifically, let Ω be a Polish space and let C be any analytic set in $\Omega \times Y$ where $Y = S \times [0,+\infty]$. Let Φ be the mapping which maps $E \subset Y$ to the set given by the expression in $\{\cdots\}$ in the definition (C.3) for Ψ, so that $\Psi(E) = E \cap \Phi(E)$. [24, Lemma 3.1] shows that Φ is uniformly analytic, i.e., the set $D = \{(\omega, y) \mid \omega \in \Omega, y \in \Phi(C_\omega)\}$ is analytic. Then, $\{(\omega, y) \mid \omega \in \Omega, y \in \Psi(C_\omega)\} = C \cap D$ is also analytic. Hence Ψ is uniformly analytic.